

# General-Equilibrium Treatment Effects: A Study of Tuition Policy

By James J Heckman, Lance Lochner,  
and Christopher Taber

*Econ 350, Winter 2014*

*The American Economic Review*, Vol. 88, No. 2, Papers and Proceedings of the Hundred and Tenth Annual Meeting of the American Economic Association (May, 1998), 381-386.

This paper considers the effects of changes in tuition on schooling and earnings, accounting for general-equilibrium effects on skill prices. The typical evaluation estimates the response of college enrollment to tuition variation using geographically dispersed cross sections of individuals facing different tuition rates. These estimates are then used to determine how subsidies to tuition will raise enrollment. The impact of tuition policies on earnings is evaluated using a schooling-earnings relationship fit on pre-intervention data and does not account for the enrollment effects of the taxes raised to finance the tuition subsidy. Thomas Kane (1994) exemplifies this approach.

The danger in this widely used practice is that what is true for policies affecting a small number of individuals need not be true for policies that affect the economy at large. A national tuition-reduction policy that stimulates substantial college enrollment will likely reduce college skill prices, as advocates of the policy claim. However, agents who account for these changes will not enroll in school at the levels calculated from conventional procedures, which ignore the impact of the induced enrollment on earnings. As a result, standard policy-evaluation practices are likely to be misleading about the effects of tuition policy on schooling attainment and wage inequality.

The empirical question is: how misleading? We show that these practices lead to estimates of enrollment responses that are more than ten times larger than the long-run general-equilibrium effects. We also improve on current practice in the treatment-effects literature by considering both the gross benefits of the program and the tax costs of financing the treatment as borne by different groups.

Evaluating the general-equilibrium effects of a national tuition policy requires more information than the tuition-enrollment parameter that is the centerpiece of partial-equilibrium policy analysis. Most policy proposals extrapolate well outside the range of known experience and ignore the effects of induced changes in skill quantities on skill prices. To improve on current practice, we have developed an empirically justified, dynamic, overlapping-generations, general-equilibrium framework for the pricing of heterogeneous skills. It is based on an empirically grounded theory of the supply of schooling and post-school human capital, where different schooling levels represent different skills.

Individuals differ in learning ability and in initial endowments of human capital. Household saving behavior generates the aggregate capital stock, and output is produced by combining the stocks of different human capitals with physical capital. The framework explains the pattern of rising wage inequality experienced in the United States in the past 30 years (Heckman et al., 1998). In this paper we apply this framework to evaluate tuition policies that attempt to increase college enrollment.

The statistical and econometric literature on “treatment effects” is remarkable for its inattention to the market consequences of the programs it evaluates. The widely used “Rubin” model (Donald Rubin, 1978) assumes no interactions among the agents being analyzed. The paradigm in the econometric literature on treatment effects is that of evaluating the effectiveness of a drug. It assumes that there are no spillovers to society at large that flow from drug use (or “treatment”) by individuals.

The literature in economics recognizes these spillover effects. The classical analysis of union relative-wage effects by H. Gregg Lewis (1963) explicitly accounts for the discrepancy between the effects of treatment (unionism) on an individual and treatment applied to an industry when prices adjust to industry-wide unionization levels. Our analysis extends Lewis's static general-equilibrium framework to a dynamic setting with skill formation.



## I. Conventional Models of Treatment Effects

The standard framework for a micro-econometric program evaluation is partial-equilibrium in character (see Heckman and Richard Robb, 1985). For a given individual  $i$ ,  $Y_{0,i}$  is defined to be the outcome the individual experiences if he does not participate in the program, and  $Y_{1,i}$  is the outcome he experiences if he does participate. The treatment effect for person  $i$  is  $\Delta_i = Y_{1,i} - Y_{0,i}$ . When interventions have general-equilibrium consequences, these effects depend on who else is treated and the market interaction between the treated and the untreated.

To see the problems that arise in the standard framework, consider instituting a national tuition policy. In this case,  $Y_{0,i}$  is person  $i$ 's wage if he does not attend college, and  $Y_{1,i}$  is his wage if he does attend. The “parameter”  $\Delta_i$  then represents the impact of college, and it can be used to estimate the impact of tuition policies on wages. It is a constant, or policy-invariant, parameter only if wages ( $Y_{0,i}$ ,  $Y_{1,i}$ ) are invariant to the number of college and high-school graduates in the economy.

In a general-equilibrium setting, an increase in tuition increases the number of individuals who attend college, which in turn decreases the relative wages of college attendees,  $Y_{1,i} / Y_{0,i}$ . In this case, the program not only impacts the wages of individuals who are induced to move by the program, but also has an impact on the wages of those who do not. For two reasons, then, the “treatment-effect” framework is inadequate. First, the parameters of interest depend on who in the economy is “treated” and who is not. Second, these parameters do not measure the full impact of the program.

For example, increasing tuition subsidies may increase the earnings of uneducated individuals who do not take advantage of the subsidy. To pay for the subsidy, the highly educated would be taxed, and this may affect their investment behavior. In addition, more competitors for educated workers enter the market as a result of the policy, and their earnings are depressed. Conventional methods ignore the effect of the policy on nonparticipants. In order to account for these effects, it is necessary to conduct a general-equilibrium analysis.

## **II. Exploring Increases in Tuition Subsidies in a General-Equilibrium Model**

We first simulate the effects of a revenue-neutral \$500 increase in tuition subsidy (financed by a proportional tax) on enrollment in college and wage inequality starting from a baseline economy that describes the United States in the mid-1980's and that produces wage growth profiles and schooling-enrollment and capital-stock data that match micro and macro statistics. The partial-equilibrium increase in college attendance is 5.3 percent. This analysis holds skill prices and, therefore, college and high-school wage rates fixed—a typical assumption in micro-economic treatment-effect analyses.

When the policy is evaluated in a general-equilibrium setting, the estimated effect falls to 0.46 percent. Because the college–high-school wage ratio falls as more individuals attend college, the returns to college are less than when the wage ratio is held fixed. Rational agents understand this effect of the tuition policy on skill prices and adjust their college-going behavior accordingly. Policy analysis of the type offered in the treatment-effect literature ignores the responses of rational agents to the policies being evaluated.

There is substantial attenuation of the effects of tuition policy on capital and the stocks of the different skills in our model. In our baseline specification, we allow skill prices and interest rates to adjust in general equilibrium but hold the pre-subsidy tuition level fixed. Simulating the policy under a number of additional alternative assumptions about the parameters of the economic model, including a case where tuition costs rise with enrollment, reproduces the basic result of substantial partial-equilibrium effects and much weaker general-equilibrium effects.

Our steady-state results are long-run effects. When we simulate the model with rational expectations, the short-run enrollment effects are also very small, as agents anticipate the effects of the policy on skill prices and calculate that there is little gain from attending college at higher rates. If we simulate using myopic expectations, the short-run enrollment effects are much closer to the estimated partial-equilibrium effects. All of these results are qualitatively robust to the choice of different tax schedules. Progressive tax schedules choke off skill investment and lead to lower enrollment responses in general equilibrium.



We next consider the impact of a policy change on discounted earnings and utility. We decompose the total effects into benefits and costs, including tax costs for each group. For the sake of brevity, we report overall results, and not the results by ability type. Table 1 compares outcomes in two steady states: (i) the benchmark steady state and (ii) the steady state associated with the new tuition policy. Given that the estimated general-equilibrium schooling response to a \$500 subsidy is small, we instead use an extremely high \$5,000 subsidy for the purpose of exploring general-equilibrium effects.

The row ‘‘high-school–high-school’’ reports the change in a variety of outcome measures for those persons who would be in high school under the benchmark or new policy regime; the high-school–college row reports the change in the same measures for high-school students in the benchmark who are induced to attend college only by the new policy; college–high-school outcomes refer to those persons in college in the benchmark economy who only attend high school after the new policy is put in place; and so forth.

By the measure of the present value of earnings, some of those induced to change are worse off. Contrary to the monotonicity assumption built into the LATE parameter of Guido Imbens and Joshua Angrist (1994), defined in this context as the effect of tuition change on the earnings of those induced to go to college, we find that the tuition policy produces a two-way flow. Some people who would have attended college in the benchmark regime no longer do so. Again, contrary to the implicit assumption built into LATE that only those who change status are affected by the policy, the rest of society also is affected by the policy.

TABLE 1—SIMULATED EFFECTS OF A \$5,000 TUITION  
 SUBSIDY ON DIFFERENT GROUPS: STEADY-STATE  
 CHANGES IN PRESENT VALUE OF LIFETIME WEALTH  
 (THOUSANDS OF 1995 DOLLARS)

Group (proportion) <sup>a</sup>	(i) After-tax earnings using base tax	(ii) After-tax earnings	(iii) After-tax earnings net of tuition	(iv) Utility
High-school– high- school (0.528)	9.512	–0.024	–0.024	–0.024
High-school– college (0.025)	–4.231	–13.446	1.529	1.411
College– high- school (0.003)	–46.711	–57.139	–53.019	–0.879
College– college (0.444)	–7.654	–18.204	0.420	0.420

*Notes:* Column (i) reports the after-tax present value of earnings in thousands of dollars discounted using the after-tax interest rate, where the tax rate used for the second steady state is the base tax rate. Column (i) reports just the effect on earnings; column (ii) adds the effect of taxes; column (iii) adds the effect of tuition subsidies; and column (iv) includes the nonpecuniary costs of college, expressed in dollars.

<sup>a</sup> The groups denote counterfactual groups. For example, the high-school–high-school group consists of individuals who would not attend college in either steady state, and the high-school–college group would not attend college in the first steady state, but would in the second, etc.

People who would have gone to college without the policy and continue to do so after the policy is put in place are financially worse off for two reasons: (i) the price of their skill is depressed, and (ii) they must pay higher taxes to finance the policy. However, they now receive a tuition subsidy, and for this reason, on net, they are slightly better off both financially and in terms of utility. Those who would abstain from attending college in both steady states are essentially indifferent between them.

They pay higher taxes, but their skill becomes more scarce, and their wages rise. Those induced to attend college by the policy are better off in terms of utility but are not better off in terms of income. Note that neither category of non-changers is a natural benchmark for a “difference in differences” estimator. The movement in their wages before and after the policy is due to the policy and cannot be attributed to a benchmark “trend” that is independent of the policy.

Table 2 presents the impact of the \$5,000 tuition policy on the log earnings of individuals with ten years of work experience for different definitions of treatment effects. The partial-equilibrium version given in the first column holds skill prices constant at initial steady-state values. The general-equilibrium version given in the second column allows prices to adjust when college enrollment varies. Consider four parameters initially defined in a partial-equilibrium context. The *average treatment effect* is defined for a randomly selected person in the population in the benchmark economy and indicates how that person would gain in wages by moving from high school to college.



The parameter *treatment on treated* is defined as the average gain over the college wage received by individuals who did not attend college. The *marginal treatment effect* is defined for individuals who are indifferent between going to college or not. It is a limit version of the LATE parameter under conventional assumptions made in discrete-choice theory (Heckman, 1997). Taber (1997) considers this parameter in his analysis of schooling choices.

Column (ii) presents the general-equilibrium version of these treatment effects. *Treatment on the treated* compares the earnings of college graduates in the benchmark economy with what they would earn if no one went to college.<sup>1</sup> *Treatment on the untreated* is defined analogously by comparing what high-school graduates in the benchmark economy would earn if everyone in the population were forced to go to college. The *average treatment effect* compares the average earnings in a world in which everyone attends college against the earnings in a world in which nobody attends college.

TABLE 2—TREATMENT-EFFECT PARAMETERS, PARTIAL-EQUILIBRIUM AND GENERAL-EQUILIBRIUM (GE) DIFFERENCE IN LOG EARNINGS: COLLEGE GRADUATES (COL) VERSUS HIGH-SCHOOL GRADUATES (HS) WITH TEN YEARS OF WORK EXPERIENCE

Parameter	(i) Prices fixed	(ii) Prices vary	(iii) Percentage of sample
Average treatment effect (ATE)	0.281	1.801	100
Treatment on treated (TT)	0.294	3.364	44.7
Treatment on untreated (TOU)	0.270	-1.225	55.3
Marginal treatment effect (MTE)	0.259	0.259	—

LATE, \$5,000 subsidy:

Partial equilibrium	0.255	—	23.6
GE (HS→COL)	0.253	0.227	2.48
LATE			
GE (COL→HS)	0.393	0.365	0.34
LATER			
GE net TLATE	—	0.244	2.82

LATE \$500 subsidy:

Partial equilibrium	0.254	—	2.37
GE (HS→COL)	0.250	0.247	0.24
LATE			
GE (COL→HS)	0.393	0.390	0.03
LATER			
GE net TLATE	—	0.264	0.27

---

*Notes:* (i) “Prices fixed” denotes the difference in log earnings between college and high-school graduates for various groups. Prices are held constant at their initial steady-state levels when wage differences are calculated. In column (ii) we allow prices to adjust in response to the change in schooling proportions when calculating wage differences. For each row, column (iii) presents the total fraction of the sample over which the parameter is defined. The LATE group denotes the effect on earnings for persons who would be induced to attend college by a tuition change. In the case of GE, LATE measures the effect on individuals induced to attend college when skill prices adjust in response to quantity movements among skill groups. The partial-equilibrium LATE measures the effect of the policy on those induced to attend college when skill prices are held constant at the benchmark level.

Such dramatic policy shifts produce large estimated effects. In contrast, the general-equilibrium marginal treatment-effect parameter considers the gain to attending college for people on the margin of indifference between attending college and attending high school. In this case, as long as the mass of people in the indifference set is negligible, the partial- and general-equilibrium parameters are the same.

The final set of parameters we consider are versions of the LATE parameter. This parameter depends on the particular intervention being studied and its magnitude. The partial-equilibrium version of LATE is defined for the outcomes of individuals induced to attend college, assuming that skill prices do not change.

The general-equilibrium version is defined for the individuals induced to attend college when prices adjust in response to the policy. The two LATE parameters are quite close to each other and are also close to the marginal treatment effect.<sup>2</sup> General-equilibrium effects change the group over which the parameter is defined compared to the partial-equilibrium case. For the \$5,000 subsidy, there are substantial price effects, and the partial-equilibrium parameter differs substantially from the general-equilibrium parameter.



We also present partial- and general-equilibrium estimates for two extensions of the LATE concept: LATER (the effect of the policy on those induced to drop out of college and go to high school; reverse LATE) and TLATE (the effect of the policy on all of those induced to change whichever direction they flow). LATER is larger than LATE, indicating that those induced to drop out of college have larger gains from dropping out than those induced to enter college have from entering. TLATE is a weighted average of LATE and LATER, with weights given by the relative proportion of people who switch in each direction.

### **III. Summary**

This paper defines and estimates general-equilibrium treatment effects. Focusing on the impact of tuition policy, we find that general-equilibrium impacts of tuition on college enrollment are an order of magnitude smaller than those reported in the literature on microeconomic treatment effects. The assumptions used to justify the LATE parameter in a microeconomic setting do not carry over to a general-equilibrium framework.

Policy changes, in general, induce two-way flows and violate the monotonicity (or one-way flow) assumption of LATE. We extend the LATE concept to allow for the two-way flows induced by the policies. We present a more comprehensive approach to program evaluation by considering both the tax and benefit consequences of the program being evaluated and placing the analysis in a market setting.