Estimating the Effect of Financial Aid Offers on College Enrollment: A Regression-Discontinuity Approach

Author(s): Wilbert van der Klaauw

Source: *International Economic Review*, Nov., 2002, Vol. 43, No. 4 (Nov., 2002), pp. 1249-1287

Published by: Wiley for the Economics Department of the University of Pennsylvania and Institute of Social and Economic Research, Osaka University

Stable URL: http://www.jstor.com/stable/826967

INTERNATIONAL ECONOMIC REVIEW
Vol. 43, No. 4, November 2002

ESTIMATING THE EFFECT OF FINANCIAL AID OFFERS ON COLLEGE ENROLLMENT: A REGRESSION–DISCONTINUITY APPROACH*

BY WILBERT VAN DER KLAAUW[1]

Department of Economics, University of North Carolina–Chapel Hill

An important problem faced by colleges and universities, that of evaluating the effect of their financial aid offers on student enrollment decisions, is complicated by the likely endogeneity of the aid offer variable in a student enrollment equation. This article shows how discontinuities in an East Coast college's aid assignment rule can be exploited to obtain credible estimates of the aid effect without having to rely on arbitrary exclusion restrictions and functional form assumptions. Semiparametric estimates based on a regression–discontinuity (RD) approach affirm the importance of financial aid as an effective instrument in competing with other colleges for students.

1. INTRODUCTION

To influence the size, quality, and composition of an incoming freshman class, a college has control over two principal instruments: the decision to offer admission to a prospective student who applies for admission and the decision of how much financial aid to offer to that student. Although the decision to deny admission is typically used to restrict entry to only those with adequate ability and promise, the financial aid decision is used to make the college more accessible to admitted students with greater financial need and to encourage a subset of those admitted, especially those with the greatest academic ability, to enroll. The specific admissions criteria and financial aid allocation mechanism chosen will reflect the college's various goals, such as maintaining or increasing the college's total enrollment, attracting higher quality students, and maintaining or improving ethnic diversity, while keeping the total costs of achieving these goals within an acceptable financial aid budget.[2] Although federal and state aid are specifically intended to favor students from lower-income families, college aid is increasingly based on

* Manuscript received November 1997; revised February 2000, November 2000.

[1] I am grateful to the C.V. Starr Center for Applied Economics at New York University for research support. I have benefited from helpful comments from Guido Imbens, Chris Flinn, Josh Angrist, Petra Todd, Jin Hahn, Ken Wolpin, Frank Vella, Robert Moffitt, Tom Kane, seminar participants at Brown, Harvard, Indiana, NYU, Princeton, Rochester, Rutgers, UNC–Chapel Hill, the 1997 summer IRP and NBER meetings, and two anonymous referees. I would like to thank Gautam Barua for excellent research assistance. Please address correspondence to: Wilbert van der Klaauw, Department of Economics, University of North Carolina–Chapel Hill, CB#3305 Gardner Hall, Chapel Hill, NC 27599-3305, USA. E-mail: *vanderkl@email.unc.edu*.

[2] For examples of theoretical and empirical analyses of admissions and financial aid offer decisions, see Barnes and Neufeld (1980), Miller (1981), Venti (1983), and Ehrenberg and Sherman (1984).

academic achievement (Carlton et al., 1995; Hoxby, 2000). A belief that the recruitment of better students will enhance the educational environment (through peer-group effects) and prestige of their institutions has led to a dramatic increase in the use of merit-based aid to compete with other colleges for the brightest students.

In determining the total financial aid budget as well as the size of individual financial aid offers, the anticipated effect of an offer on a student's decision to enroll in the college plays a crucial role.[3] Although a college's past records of financial aid offers and student enrollment decisions provide a valuable source of data and experience in this regard, the actual evaluation of the effect of financial aid offers on student enrollment decisions remains a complicated matter.

A student's decision whether or not to enroll in a particular college is influenced by a number of different factors, many of which are unobserved by college administrators. The most important piece of information that is typically missing is information on a student's alternative options. These options may include admission and financial aid offers from other colleges and the option to join the labor force directly after high school. Although student application and financial aid request forms often provide some information concerning other colleges and universities the student has applied to, it is generally not known whether these colleges will admit the student and what their financial aid offers will be.[4] In addition, typically little or nothing is known about any employment options each applicant may have. This lack of information not only pertains to new applicants, but also to applicants in previous years. Most colleges do not collect information about the alternative options of those who enrolled, and about the destinations of applicants who chose not to enroll. In particular, for the latter group, it is not known whether they decided to go to college at all, and if so, in which other colleges these students chose to enroll. Colleges that do collect information about this group typically do so only for the students to whom they offered aid and collect little information about the students' options beyond the name of the college they actually enrolled in.

This lack of information about alternative options makes it very difficult to distinguish the effect of a college's own offers of financial aid from the (unobserved) offers of other colleges as well as possible alternative opportunities in the labor market. In general, we could expect each college's aid offer to depend (at least in part) on the same student characteristics, such as available measures of academic ability, ethnicity, and family income, typically reported on federal and college aid application forms. Missing information about alternative opportunities is therefore likely to cause an omitted variable bias in estimating the effect of financial aid on enrollment.

The evaluation problem is further complicated by the fact that financial aid decisions can rarely be described completely in terms of measured student characteristics. To some extent the financial aid decision is a subjective one, depending

[3] Depending on the magnitude of the aid effect and on a college's objectives, it may, for example, be optimal for a college to offer more aid to a weaker student if the probability of getting a better student is very low and aid allocation is constrained by a fixed budget.

[4] It is illegal for colleges to share information on individual students' aid offers (Carlton et al., 1995).

on an admission officer's assessment of the student's complete "package," of which certain aspects, such as statements of purpose, extracurricular activities, and recommendation letters, are typically not kept in a computer database. It is possible that, even if the student's outside opportunities were known and controlled for, some of these unmeasured aspects may be correlated with the remaining error term in an enrollment equation. For example, in choosing between two different colleges, students who are offered more aid because of their athletic talent may, all else being equal, have a preference for the college with the highest reputation in sports. The resulting dissimilarity in characteristics (academic talent, in this case), observed by the financial aid officer but not by the econometrician, between individuals receiving different amounts of financial aid leads to a second omitted variable or selection bias problem.

Because of both omitted variable problems, when evaluating the effect of financial aid on enrollment, the former cannot be considered exogenous with respect to the enrollment decision. For example, when comparing the enrollment rates of two groups of applicants who differ in the amount of financial aid they were offered but are equal in all measured characteristics, it is quite possible that the enrollment rate of the group who received more aid was actually lower. This would be the case, if those who received more had unmeasured (but observed by college financial aid officers) characteristics, such as special awards, recommendation letters, and extracurricular activities, which made them likely to have received similar or possibly better aid offers from other, perhaps more attractive colleges.

In their survey of empirical studies on the demand for higher education, Leslie and Brinkman (1988) found considerable variation in reported tuition and financial aid effect estimates. As has also been the case for more recent studies, this variation has been especially large among studies that use cross-sectional data on individual enrollment decisions from different groups of colleges and universities, or from a single institution. Although most of these studies report a positive and statistically significant effect of financial aid on enrollment, the implied magnitude of the impact varies from negligible or small (Seneca and Taussig, 1987; Parker and Summers, 1993) to considerable (Ehrenberg and Sherman, 1984; Moore et al., 1991).

To some extent, one may expect the effect estimates to differ for alternative samples of institutions and students. The enrollment effect of an increase in financial aid is likely to vary with each institution's tuition level and financial aid policy, the characteristics of its students and its applicant pool, as well as its competitors. For example, several studies (Leslie and Brinkman, 1988; Schwartz, 1985, 1986; St. John, 1990; McPherson and Schapiro, 1991) have found high school graduates from higher-income families to be less sensitive to financial aid offers in their college enrollment decisions than graduates from lower-income families.

However, especially given the wide range of the estimates and the potential for endogeneity bias discussed earlier, it is likely that differences across studies in the econometric specification of the enrollment equation and in the estimation method used contributed significantly to the variability in the reported effect estimates. It is standard practice to treat the aid award as an exogenous explanatory variable in the enrollment equation, with the implicit assumption that other included explanatory variables adequately control for a potential endogeneity bias. Given the large variation across studies in the number and types of additional regressor variables included, and in light of the omitted variable problems discussed earlier, we can therefore expect a fair amount of heterogeneity in the size and reliability of the reported effect estimates due to remaining biases. Some of the institutionally based studies of individual enrollment decisions (such as Ehrenberg and Sherman, Seneca and Taussig, and Moore et al.) incorporated specially collected information about the alternative college each student most likely would have attended or actually did enroll in. Although including the information about alternative aid offers is likely to reduce the potential for endogeneity bias, as discussed earlier, there is no guarantee that the college aid amount will be uncorrelated with other omitted variables influencing the individual's enrollment decision.

In this article, I analyze the effect of financial aid offers on the enrollment decisions of a large sample of individuals admitted to an East Coast college, referred to as College X, during the period from 1989 to 1993. This database has the usual shortcomings discussed above, in that it lacks important information on each student's choice alternatives as well as some information (such as statements of purpose, reference letters, and characteristics of the financial aid officer) that influenced the financial aid decision, but that was not included in the database. I will show, however, that it is possible to exploit idiosyncratic features of the financial aid decision process to obtain a reliable estimate of the effect of financial aid offers on enrollment. More precisely, a key component of the aid allocation decision is a simple formula used to divide students into a few groups on the basis of a calculated continuous measure of academic ability. Based on the particular interval this ability measure fell into, each applicant was assigned a rank that in turn was used to determine the aid offer. As a result, the assignment rule has features (discontinuities) that make it similar to that of a powerful, but relatively ignored quasi-experimental design: the regression–discontinuity (RD) design originally introduced by Thistlethwaite and Campbell (1960). I will show how this design can be used to obtain credible estimates of the financial aid effect without having to rely, as in commonly used selection bias correction methods, on arbitrary exclusion restrictions and functional form and distributional assumptions on errors. The enrollment elasticity estimates affirm the importance of financial aid as an effective tool to compete with other colleges for students.

More generally, this study provides an illustration of how knowledge of aspects of the selection mechanism can aid in obtaining reliable program effect estimates. Moreover, it is argued that similar features, which characterize the so-called *fuzzy RD design*, should be relatively easy to incorporate and are already likely to be found in the assignment or selection rules of many existing nonexperimental evaluation designs. In fact, it will be shown that, without making an explicit connection to the design, several recent studies, such as Angrist and Krueger's (1991) well-known evaluation of the returns to schooling, rely heavily on the RD design in identifying causal effects.

In the next section, I present a simple model of student enrollment decisions and discuss the particular features of the financial aid allocation process at College X that will be exploited to obtain estimates of the effect of financial aid on enrollment. Section 3 considers the identification and estimation of causal effects in case of a RD design and discusses how, in case of College X, the design can be used to estimate the financial aid effect on enrollment. Section 4 discusses the dataset, and Section 5 presents estimates and results from a sensitivity analysis. Finally, Section 6 provides a conclusion and discusses the applicability of the RD approach in economic evaluation studies.

2. STUDENT ENROLLMENT DECISIONS AND THE FINANCIAL AID ALLOCATION PROCESS

As outlined in Manski and Wise (1983, Chap. 2), the complete admission process can be treated as a sequence of decisions involving different agents. The first stage represents the student's college application decision. The second is the institution's admission decision. The third is the institution's financial aid offer decision, conditional on an offer of admission, and the fourth is the student's enrollment decision. In this section, I will present a simple behavioral model of stage four and provide an econometric specification of the financial aid rule describing stage three.

The student's decision problem can be characterized as having to make an optimal choice from a discrete set of school and nonschool characteristics. In our case, where we model the choice to enroll at College X, we can define a student's options, given his prior college application decisions and each college's subsequent admission and financial aid decisions, as (i) enroll at College X (and accept its financial aid offer) and (ii) enroll at another college.[5] The enrollment decision can be thought of as involving a comparison between the student's utility associated with each choice alternative. The utility a student receives from each decision will depend on the total costs and total benefits associated with each choice. The costs include tuition and living expenses minus financial aid, whereas the benefits include the consumption value of a college education as well as the student's expected future earnings and job prospects after graduating from college.

With missing data on tuition cost, living expenses, and post-graduation earnings for different colleges, I will specify the utility associated with each choice alternative simply as a linear function of the financial aid amount offered and an unobserved component capturing all other factors. Let F represent the amount of discretionary aid offered by College X and F^o the financial aid offer made by the most preferred college other than College X. Then, for an individual i the difference in utility associated with the choice to enroll or not can be defined as

$$EN_i^* = \delta(F_i - F_i^o) + v_i \tag{1}$$

where the unobserved random component v_i measures all other individual differences in utility associated with alternative choice options. Heterogeneity in the sensitivity to financial aid can further be captured by an individual specific δ_i, which varies across individuals.

[5] I ignore the option to join the labor market or military. Few of those admitted to College X tend to choose this option.

The enrollment decision therefore depends on the amount of discretionary aid offered by College X, as well as the financial aid offer made by the most preferred alternative college. With F_i^o unobserved, the utility difference can be written as

$$EN_i^* = \delta F_i + u_i \tag{2}$$

where $u_i = v_i - \delta F_i^o$.[6] Generally, one would expect financial aid offers from other colleges, and thus that of the student's most preferred alternative college option, to be correlated with College X's aid offer F_i, as they will depend on similar sets of student characteristics. In addition, F_i may be correlated with the unobserved preference component v_i. For example, students with great athletic talent may be offered more aid by College X, but may also be more or less attracted to College X, irrespective of the aid offer, because of its reputation in sports. For these reasons, u_i and F_i are likely to be correlated.

With $EN_i = 1$ if $EN_i^* > 0$ and $EN_i = 0$ otherwise, the probability that the student will enroll at College X is given by

$$\begin{aligned} \Pr(EN_i = 1) &= \Pr(\delta F_i + u_i > 0) \\ \Pr(EN_i = 0) &= 1 - \Pr(EN_i = 1) \end{aligned} \tag{3}$$

With a normally distributed unobservable the enrollment decision would be described by a Probit model. Alternatively, we can consider a linear probability model specification for (3):

$$EN_i = \beta + \alpha F_i + w_i \tag{4}$$

where w_i, like u_i, is expected to be correlated with F_i.[7]

Solving the endogeneity problem requires additional knowledge or assumptions about the financial aid assignment rule. Commonly, in estimating treatment effects more generally, to solve this type of endogeneity problem, assumptions are made that take the form of exclusion restrictions (variables that are assumed to influence treatment selection but not outcomes), index assumptions on the selection process, and/or functional form and distributional assumptions about error distributions, such as the conditional (on the index or propensity score) independence assumption adopted by matching methods.[8] In practice, few if any of these

[6] Instead of motivating the enrollment Equation (2) as the second stage of a two-stage choice problem, one could alternatively view the equation as a linear approximation to a multinomial choice problem where the error term contains the financial aid awards and characteristics of all colleges considered in the choice decision.

[7] To see this, assume that u_i can be decomposed into $u_i = u_{1i} + u_{2i}$, where u_{2i} is a uniform random variable that is independently distributed from u_{1i} and F_i. Then we can write $\Pr(EN_i = 1 | F_i, u_{1i}) = \beta + (1/\theta)(\delta F_i + u_{1i})$ with $\theta > 0$, leading to the regression equation (4) with $\alpha = \delta/\theta$, and where $E[w_i | F_i] = (1/\theta)E[u_{1i} \mid F_i]$. See Heckman and Snyder (1997) for a general derivation of the linear probability model as a representation of a random utility model.

[8] For discussions of such approaches and the identifiability of treatment effects, see Rosenbaum and Rubin (1983), Heckman and Robb (1985), Heckman and Hotz (1989), Heckman (1990), Manski (1990), Imbens and Angrist (1994), and Heckman et al. (1997).

assumptions about the selection rule are based on actual knowledge of the selection or assignment procedure. In the case considered here, however, where we want to estimate the effect of financial aid offers on the enrollment decisions of students admitted to College X, there is additional information about the college's financial aid rule that can be exploited to solve this endogeneity bias problem and to obtain credible estimates.

The actual decision rule adopted by College X in determining each student's financial aid offer is fairly complex, involves both objective and subjective evaluations, and is therefore difficult to characterize by a simple formula. In addition, although many relevant student characteristics, such as various academic ability measures, minority status, and (for some) parental income are included in the college's database, others, such as the student's statement of purpose, extracurricular activities, transcripts, and recommendation letters, are not. An important feature of the financial aid decision process of most colleges, however, is the existence of simple rules, the specifics of which are generally unknown to student applicants, which are designed to make the allocation of aid more objective, regulated, and easier to implement. One such rule, adopted by College X, is the use of a simple formula that converts a student's SAT scores and high school grade point average (GPA) into an index that is then used to rank students into a small number of categories. More precisely, during the period studied, the particular index formula used was

$$S = \phi_0 \times (\text{first three digits of total SAT score}) + \phi_1 \times \text{GPA}$$

where ϕ_0 and ϕ_1 are known weights and S represents the calculated index. Applicants were then divided into four groups on the basis of the interval the calculated index fell into. These intervals were determined by three cutoff points on the S scale. Let the three cutoff points in ascending order be denoted by $\bar{S}_1$, $\bar{S}_2$, and $\bar{S}_3$, respectively; then the highest rank or category would consist of students with index scores above $\bar{S}_3$.

Students of different rank are eligible for different amounts of aid. Within a rank, a base amount is assigned that is subsequently adjusted on the basis of a student's minority status, family income, as well as more subjective and detailed evaluations of the strength of the student's complete application package. These adjustments can therefore be merit, affirmative action, or need based. It is possible that these adjustments may to some extent depend on, or be correlated with, the value of the individual's GPA and SAT scores and thus with the ability index S. Therefore, within each rank group, it may be the case that those with a higher index score receive on average somewhat more or somewhat less aid. However, given the predominant importance of each individual's rank in the aid allocation process (as will be shown later on), the average aid offer as a function of S will contain jumps at the cutoff points for the different ranks, with those scoring just below a cutoff point receiving considerably less on average than individuals who scored just above the cutoff point.

With F_i representing the total amount of discretionary college aid offered by College X to individual i, the financial aid allocation process just described can be characterized as

(5)

$$F_i = E[F_i \mid S_i] + e_i = f(S_i) + \gamma_1 \cdot 1\{S_i \geq \bar{S}_1\} + \gamma_2 \cdot 1\{S_i \geq \bar{S}_2\} + \gamma_3 \cdot 1\{S_i \geq \bar{S}_3\} + e_i$$

where $f(S)$ is some continuous function of S, $1\{\}$ is the indicator function, taking the value one if the logical condition in brackets holds and the value zero if not, and e_i is an unobserved component, capturing the residual effect of all other relevant characteristics of the student (and possibly of the financial aid officer) influencing the financial aid decision.

Because of the discontinuities in the average amount of financial aid offered as a function of the composite ability measure, the assignment mechanism conforms to that of the RD design. This quasi-experimental evaluation design was first introduced by Thistlethwaite and Campbell (1960). Although the design has been frequently discussed since then as a potentially powerful quasi-experimental design in the evaluation literature (see, for example, Campbell and Stanley, 1963; Cook and Campbell, 1979; Trochim, 1984), until recently the design had been largely ignored by economists with the notable exception of some early papers by Goldberger (1972), Cain (1975), and Barnow et al. (1980).[9] Moreover, in spite of a growing number of applications of the RD approach, including two recent applications by Angrist and Lavy (1999) and Pitt and Khandker (1998) in economics, the program evaluation literature has long lacked a formal discussion of several important issues relating to the identifiability of treatment effects and to the development of alternative estimation methods for the RD design.[10] This empirical study and the theoretical article by Hahn et al. (2001) address this gap in the literature.

To see how knowledge of discontinuities in the financial aid assignment rule can be exploited to estimate causal effects, it is useful to first consider a simpler evaluation case with a single treatment dose level and where the assignment rule has a single discontinuity. This will be followed by a consideration of the case with multiple discontinuities and multiple treatment dose levels, which applies to our evaluation problem of estimating the effect of financial aid awards on enrollment decisions. The discussion will be closely related to that in Hahn, Todd, and Van der Klaauw (HTV in what follows), but differs from it in several aspects. First, I

[9] Given the lack of attention it received in applied economic research, the design was only briefly mentioned in Meyer's (1995) survey of the different evaluation approaches used in economics.

[10] Among others, the design has been applied to evaluate the impact of a National Merit Award on obtaining additional college scholarships and on career aspirations (Thistlethwaite and Campbell, 1960), the effect on subsequent student performance of being placed on the dean's list (Seaver and Quarton, 1976), the effect of various compensatory education programs (Trochim, 1984), the effect of extending unemployment benefits to released prisoners on recidivism rates (Berk and Rauma, 1983), the effect of a program designed to educate employees about their lifestyles in relation to risk of heart disease (Visser and de Leeuw, 1984), the effect of the Research Career Development Award program of the NIH on research productivity (Carter et al., 1987), and the impact of an accelerated mathematics program for gifted children (Robinson and Stanley, 1989). In the recent economic application by Angrist and Lavy, an RD-motivated IV approach is used to estimate the effect of class size on student performance in Israel, whereas Pitt and Khandker use FIML methods to estimate the impact on household behavior of a group-based credit program in Bangladesh.

explore the issues of identification and estimation within the context of the more traditional econometric regression model for program evaluation and consider the case of single, as well as multiple treatment dose levels and discontinuities. Second, the discussion is used to motivate and to help assess the reliability of an alternative two-stage semiparametric estimation approach.

3. THE REGRESSION–DISCONTINUITY APPROACH

Consider a random sample of individuals where for each individual i we observe an outcome measure Y_i (for example, the individual's enrollment decision) and a treatment indicator T_i, equal to one if treatment was received and zero otherwise (for example, an offer of a fixed scholarship amount). A common regression model representation of the evaluation problem is

$$Y_i = \beta + \alpha \cdot T_i + u_i \tag{6}$$

In case of a constant treatment effect, when assignment to (or self-selection into) treatment is nonrandom, selection bias in the estimation of α can arise because of a dependence between T_i and u_i. In this case, $E[u \mid T] \neq 0$ and the endogeneity of T will generally render the OLS estimate of α inconsistent. Similarly, in the case where the treatment effect varies across individuals, the OLS estimate of the treatment variable coefficient will generally not have a causal interpretation, whereas in the case of randomized assignment, it would estimate $E[\alpha_i]$, the average treatment effect in the population.

In case of a RD design, we have additional information about the selection rule: It is known that the treatment assignment mechanism depends (at least in part) on the value of an observed continuous variable relative to a given cutoff score, in such a way that the corresponding propensity score (the probability of receiving treatment) is a discontinuous function of this variable at that cutoff score. In the simplest and most frequently discussed version of the design, referred to in the literature as the *sharp RD design* (Trochim, 1984), individuals are assigned to treatment and control groups solely on the basis of an observed continuous measure S, called the *selection* or *assignment variable*. Those who fall below some distinct cutoff point $\bar{S}$ are placed in the control group ($T_i = 0$), whereas those on or above that point are placed in the treatment group ($T_i = 1$) (or vice versa). Thus, assignment occurs through a known and measured deterministic decision rule: $T_i = T(S_i) = 1\{S_i \geq \bar{S}\}$. The variable S itself may well be directly related to the outcome Y. That would automatically cause T to be related to Y as well, even if the treatment had no causal effect on Y. This is in sharp contrast with pure randomization. Although randomization guarantees that treatment and control groups will be as similar as possible in characteristics other than the treatment itself, the sharp RD design makes them very different, at least in terms of their average S value. The design also violates the "strong ignorability condition" of Rosenbaum and Rubin (1983), which, in addition to requiring u to be independent of T conditional on S, requires $0 < \Pr(T = 1 \mid S) < 1$ for all S, whereas here $\Pr(T = 1 \mid S) \in \{0, 1\}$.[11]

[11] In the terminology of Heckman et al. (1997), there is no region of common support.

To see how treatment effects can be identified and estimated in case of the RD design, note that consideration of the sample of individuals within a very small interval around the cutoff point will be very similar to a randomized experiment at the cutoff point (a tie-breaking experiment). That is, because they have essentially the same S value, we can expect individuals just below the cutoff score on average to be very similar to individuals just above the cutoff point and thus to have similar average outcomes in the absence of the program as well as similar average outcomes when receiving treatment. With those to the right of the cutoff receiving treatment and those to the left not, a comparison of the average outcomes of both groups should therefore provide a good estimate of the treatment effect. Of course, in the case of varying treatment effects, the estimate will only apply to the subset of individuals close to the cutoff point.

Increasing the interval around the cutoff point is likely to produce a bias in the effect estimate, especially if the assignment variable was itself related to the outcome variable conditional on treatment status. If an assumption is made about the functional form of this relationship between the average outcome and the selection variable, on the other hand, we can use more observations and extrapolate from above and below the cutoff point to what a tie-breaking randomized experiment would have shown. This double extrapolation combined with the exploitation of the "randomized experiment" around the cutoff point has been the main idea behind regression–discontinuity analysis.

We can analyze the identifiability of a constant treatment effect in this RD design more formally by observing that $\lim_{S\downarrow\bar{S}} E[Y \mid S] - \lim_{S\uparrow\bar{S}} E[Y \mid S] = \alpha + \lim_{S\downarrow\bar{S}} E[u \mid S] - \lim_{S\uparrow\bar{S}} E[u \mid S]$. The following definition represents a more formal way of assuming that in the absence of treatment, individuals in a small interval around $\bar{S}$ would have similar average outcomes.

Assumption A1. *The conditional mean function $E[u \mid S]$ is continuous at $\bar{S}$.*

Then it follows that under assumption A1 the treatment effect α will be identified by the difference[12]

$$\lim_{S\downarrow\bar{S}} E[Y \mid S] - \lim_{S\uparrow\bar{S}} E[Y \mid S] \tag{7}$$

For the case of varying treatment effects, where the treatment effect may be a function of S or vary randomly across individuals, we introduce the following continuity assumption.

Assumption A2. *The mean treatment effect function $E[\alpha_i \mid S]$ is right-continuous at $\bar{S}$.*

Then, under assumptions A1 and A2, the difference defined in (7) will equal $E[\alpha_i \mid S = \bar{S}]$, the average treatment effect of those at the margin. In the case

[12] I implicitly assume that the density of S is positive in the area around $\bar{S}$ and that both limits exist.

where treatment effects vary in a deterministic way with S as in $Y_i = \beta + \alpha(S_i) \cdot T_i + u_i$, then under assumptions A1 and A2 (requiring the nonparametric function $\alpha(S)$ to be right-continuous at $\bar{S}$), the difference defined in (7) will identify the local treatment effect $\alpha(\bar{S})$ at $\bar{S}$.

Considering the estimation of a constant treatment effect with a sharp RD design, note that this design is a special case of selection on observables, that is, where the dependence between T_i and u_i arises because the treatment status of an individual is related to some observed characteristic that itself is related to the outcome variable. With $T_i = T(S_i) = 1\{S_i \geq \bar{S}\}$, a dependence between the assignment variable S_i and the error term u_i would generally lead to a biased and inconsistent OLS estimate of α. One approach to estimate the treatment effect in this case is to specify and include the conditional mean function $E[u \mid T, S]$ as a "control function" in the outcome equation (Heckman and Robb, 1985).

In the sharp RD design case, $E[u \mid T, S] = E[u \mid S]$, that is, since S is the only systematic determinant of treatment status T, S will capture any correlation between T and u. As a result, by entering the correct specification of the control function $k(S)$ alongside T, the equation

$$Y_i = \beta + \alpha \cdot T_i + k(S_i) + \omega_i \tag{8}$$

where $\omega_i = Y_i - E[Y_i \mid T_i, S_i]$, can be estimated to yield a consistent estimate of the program effect, as it will free T from the contamination that leads to a selection bias.[13] Goldberger (1972) and Cain (1975) considered the case of a linear control function. As illustrated in Figure 1, α will then be estimated by the distance between the two linear, parallel regression lines at the cutoff point, which in this case equals the difference in the intercepts of the two lines. This will be an unbiased estimate of the common treatment effect if the control function is correctly specified, that is, if the true conditional mean function $E[u \mid S]$ is in fact linear.

Similarly, in case of treatment effects that vary deterministically with S, we can obtain an estimate of the local treatment effect $\alpha(\bar{S})$ by estimating the control function-augmented regression function (8) but now with $\alpha(\bar{S})$ replacing α and the control function $k(S)$ representing a specification of the function $E[u_i \mid S] + (\alpha(S) - \alpha(\bar{S}))1\{S \geq \bar{S}\}$. In the more general case of varying treatment effects, we can estimate a version of Equation (8) in which the coefficient on T equals the local average treatment effect $E[\alpha_i \mid \bar{S}]$ and where $k(S)$ is a specification of the function $E[u_i \mid S] + (E[\alpha_i \mid S] - E[\alpha_i \mid \bar{S}])1\{S \geq \bar{S}\}$.[14]

This regression-based estimation approach requires a specification $k(S)$ of the control function. Although a misspecified control function is likely to produce inconsistent estimates, the identification conditions A1 and A2 only impose that

[13] This point that unbiasedness is attainable when the variables that determine the assignment are known, quantified, and included in the equation was originally made by Goldberger (1972). His analysis, however, presumes that the relationship between these variables and the outcome variable is known to be linear.

[14] Note that if data were also available from a setting in which no program was present, one would actually be able to uncover $E[u_i \mid S]$, in which case we could identify $E[\alpha_i \mid S]$ for all $S \geq \bar{S}$.

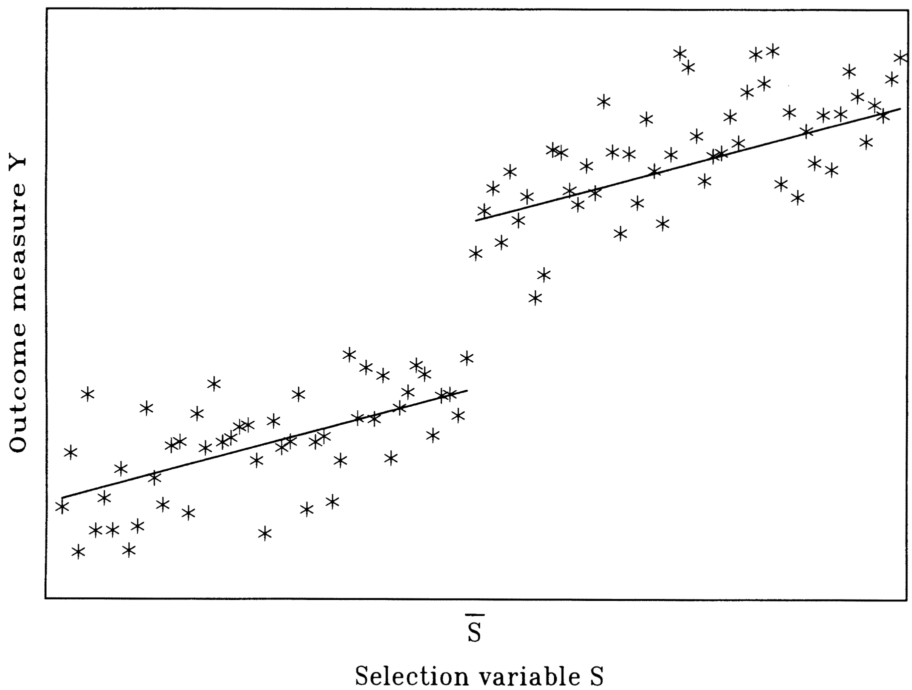

FIGURE 1

REGRESSION–DISCONTINUITY DATA

the control and average treatment effect functions be continuous in S at the cut-off point in each of the cases discussed above. To minimize the potential for misspecification, one could therefore adopt a semiparametric specification for the control function, or use local or nonparametric regression around the cut-off point $\bar{S}$. As long as the control function is continuous in S, identification will be guaranteed because of the discontinuity in the function $T(S)$. In this article, I will use a semiparametric approach, similar in spirit to that proposed by Trochim (1984). It uses a power series approximation for $k(S) \approx \sum_{j=1}^{J} \eta_j \cdot S^j$, where the number of power functions, J, is estimated from the data by generalized cross-validation (see, for example, Newey et al., 1990).

An alternative estimation approach would be to estimate the limits $\lim_{S\downarrow\bar{S}} E[Y \mid S]$ and $\lim_{S\uparrow\bar{S}} E[Y \mid S]$ nonparametrically. The difference between two consistent estimators of these limits would then be a consistent estimator of the (local) treatment effect. Two such estimators were proposed by HTV: a local linear regression estimator and a Kernel regression estimator.

3.1. *The Fuzzy RD design.* Although in the sharp RD design, treatment assignment is known to depend on the selection variable S in a deterministic way, in the second type of the RD design, referred to in the literature as the *fuzzy RD design* (Campbell, 1969), treatment assignment depends on S in a stochastic

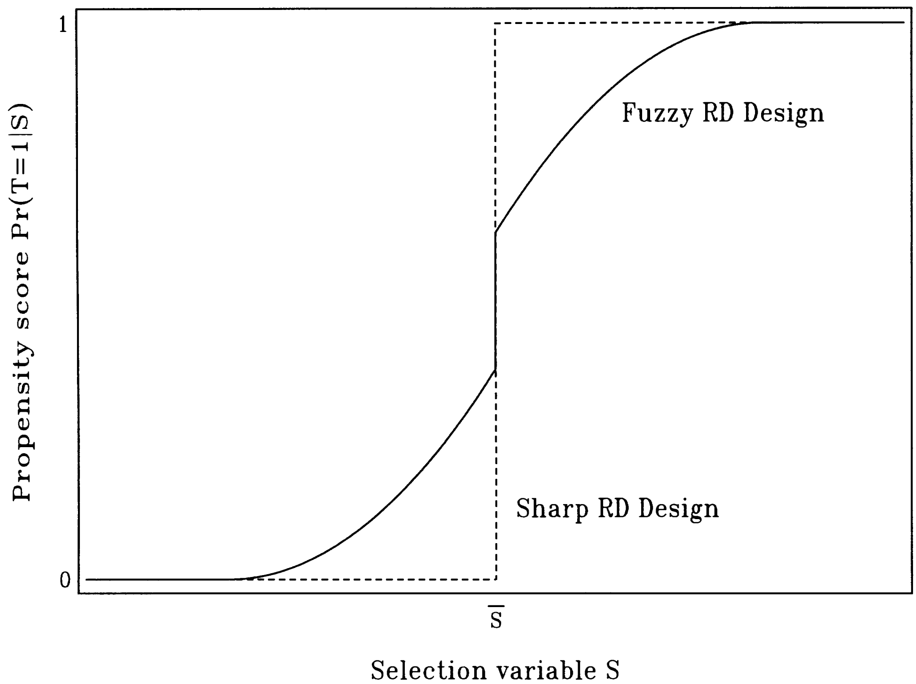

FIGURE 2

ASSIGNMENT IN THE SHARP (DASHED) AND FUZZY (SOLID) RD DESIGN

manner but in such a way that the propensity score function $\Pr(T = 1 \mid S)$ is again known to have a discontinuity at $\bar{S}$. The fuzzy design can occur in case of misassignment relative to the cutoff value in a sharp design, with values of S near the cutoff appearing in both treatment and control groups. Alternatively, in addition to the position of the individual's score relative to the cutoff value, assignment may be based on additional variables observed by the administrator, but unobserved by the evaluator. Compared to the sharp design, selection here is both on observables and unobservables. Instead of having the step function $\Pr(T = 1 \mid S) = 1\{S \geq \bar{S}\}$, the selection probability as a function of S may now appear as the S-shaped function shown in Figure 2.

As in the sharp RD design case, it is again possible to exploit the discontinuity in the selection rule to identify a treatment effect under continuity assumption A1. To see this, note that if the conditional mean function $E[u \mid S]$ is continuous at $S = \bar{S}$, then $\lim_{S \downarrow \bar{S}} E[Y \mid S] - \lim_{S \uparrow \bar{S}} E[Y \mid S] = \alpha(\lim_{S \downarrow \bar{S}} E[T \mid S] - \lim_{S \uparrow \bar{S}} E[T \mid S])$. It follows that the treatment effect α is identified by

$$(9) \qquad \frac{\lim_{S \downarrow \bar{S}} E[Y \mid S] - \lim_{S \uparrow \bar{S}} E[Y \mid S]}{\lim_{S \downarrow \bar{S}} E[T \mid S] - \lim_{S \uparrow \bar{S}} E[T \mid S]}$$

where the denominator in (9) is nonzero because of the known discontinuity of $E[T \mid S]$ at $\bar{S}$.

To analyze the identifiability of treatment effects in the case of heterogeneous or varying treatment responses, consider the following two assumptions.

ASSUMPTION A2*. *The average treatment effect function* $E[\alpha_i \mid S]$ *is continuous at* $\bar{S}$.

ASSUMPTION A3. T_i *is independent of* α_i *conditional on* S *near* $\bar{S}$.

ASSUMPTION A2* is a generalization of assumption A2, and is a formal way of assuming that the average treatment effect is similar for individuals with values of S close to $\bar{S}$. When α is a deterministic function of S, then under assumptions A1 and A2*, the ratio in (9) identifies the local treatment effect $\alpha(\bar{S})$. Similarly, in the more general case of varying treatment effects, it is straightforward to show that, under continuity assumptions A1 and A2* and the local conditional independence assumption A3, (9) identifies $E[\alpha_i \mid S = \bar{S}]$, the average treatment effect at the cutoff point.[15]

In general, as discussed by Barnow et al. (1980), if the selection process is not perfectly known and the unknown component e_i in $T_i = E[T_i \mid S_i] + e_i$ (where $E[T_i \mid S_i] = \Pr[T_i = 1 \mid S_i]$) is not an independent assignment error (that is, independent of u_i given S_i), estimation of the control function-augmented outcome equation (8) will lead to biased estimates.[16] The bias will depend on the covariance of T and u conditional on S and may be positive or negative. We can solve this selection bias problem and obtain an estimate of the treatment effect α, by estimating the same control function-augmented outcome equation, but where T_i is now replaced by the propensity score $E[T_i \mid S_i]$,

$$Y_i = \beta + \alpha \cdot E[T_i \mid S_i] + k(S_i) + w_i \tag{10}$$

where $w_i = Y_i - E[Y_i \mid S_i]$ and $k(S_i)$ is a specification of $E[u_i \mid S_i]$. This suggests a two-stage procedure, where, in the first stage, we specify the propensity score function in the fuzzy RD design as

$$E[T_i \mid S_i] = f(S_i) + \gamma \cdot 1\{S_i \geq \bar{S}\} \tag{11}$$

where $f(\cdot)$ is some function of S that is continuous at $\bar{S}$. By specifying the functional form of f (or by estimating f semi- or nonparametrically) we can estimate

[15] The conditional independence assumption is a strong assumption that may be violated if individuals self-select into or are selected for treatment on the basis of expected gains from treatment. HTV show that under a less restrictive local monotonicity assumption similar to that proposed by Imbens and Angrist (1994), the ratio defined in (9) identifies a local average treatment effect at the cutoff point, for the subgroup of individuals for whom treatment changes discontinuously at the cutoff point.

[16] In the "mixed model" of Cain (1975, p. 309), assignment is based on the selection variable and an independent assignment error e, in which case, after controlling for selection on observables through the control function $k(S)$ (or through matching methods), there will be no remaining selection bias in the treatment effect estimate. Note that, in this case, it would be possible to identify α nonparametrically, and, in the case of varying treatment effects, one could nonparametrically identify the average treatment effect over the region of common support.

γ, the discontinuity in the propensity score function at $\bar{S}$. In the second stage, the control function-augmented outcome equation is then estimated with T_i replaced by the first-stage estimate of $E[T_i \mid S_i] = \Pr[T_i = 1 \mid S_i]$ as in Maddala and Lee (1976).[17] Note that this estimate will be discontinuous in S (unless $\hat{\gamma} = 0$), whereas the included control function $k(S)$ will be continuous in S at $\bar{S}$. Under the correct specifications for $k(S)$ and $f(S)$, this two-stage procedure will produce a consistent estimate of the treatment effect.[18]

When treatment effects vary deterministically with S, a two-stage estimate of $\alpha(\bar{S})$ can be obtained by estimating the control function-augmented regression function (10) but now with $\alpha(\bar{S})$ replacing α and $k(S)$ representing a specification of the function $E[u_i \mid S] + (\alpha(S) - \alpha(\bar{S}))E[T \mid S]$. Note again that under assumptions A1 and A2*, this function will be continuous at $\bar{S}$. Similarly, in the case of varying treatment effects, under assumptions A1, A2*, and A3, we can estimate a version of Equation (10) in which α is replaced by $E[\alpha_i \mid \bar{S}]$ and where $k(S)$ is a specification of the function $E[u_i \mid S] + (E[\alpha_i \mid S] - E[\alpha_i \mid \bar{S}])E[T \mid S]$, which again will be continuous at $\bar{S}$.

An alternative, nonparametric approach proposed by HTV is to divide a nonparametric estimator of the difference in (7) discussed earlier for the sharp design, by a nonparametric estimator of the difference between $\lim_{S \downarrow \bar{S}} E[T \mid S]$ and $\lim_{S \uparrow \bar{S}} E[T \mid S]$ to obtain a consistent estimator of the treatment effect identified by (9). One such estimator, which is based on one-sided Kernel regression, was shown by HTV to be numerically equivalent to a local Wald estimator. I will discuss and apply this estimator later in this article and compare the resulting estimate with that obtained using the two-stage estimation approach described above.

In the evaluation problem considered in this article, which is to estimate the effect of financial aid offers on enrollment decisions, the treatment variable is F_i, the amount of college aid offered, which is not binary, but takes on many different values, and the average treatment level as a function of the ability index S, $E[F_i \mid S]$, contains three discontinuities. The outcome variable in our case is the enrollment decision EN_i, and we can interpret the enrollment equation (4), where $EN_i = \beta + \alpha F_i + w_i$, as describing the average potential enrollment outcomes across individuals under alternative financial aid award assignments. With a common linear financial aid effect all three discontinuities help identify α. As long as the conditional expectation function $E[w_i \mid S]$ is continuous in S at each discontinuity point $\bar{S}_1$, $\bar{S}_2$, and $\bar{S}_3$, each of the corresponding ratios defined in (9) identifies the aid effect.

[17] In the RD design literature, Spiegelman (1979) and Trochim and Spiegelman (1980) proposed a similar two-stage method, but where, in the first stage, the propensity score function was estimated using a nearest-neighbor moving average method, which involved computing the moving average of the T values for cases ordered by S. Note that this approach produces an estimate of $E[T \mid S]$ that is continuous in S, implying that, under our identification conditions, a treatment effect would no longer be identified.

[18] Note that in the case of a parametric approach, if we assume the same functional form for $k(S)$ as for the function $f(S)$ in the first-stage regression $T_i = f(S_i) + \gamma \cdot 1\{S_i \geq \bar{S}\} + e_i$, then the two-stage estimation procedure described here will be equivalent to two-stage least squares (in the case of a linear-in-parameter specification) with $1\{S_i \geq \bar{S}\}$ and the terms in $f(S)$ as instruments.

We may allow the financial aid effect to vary with the ability index S. In that case, under the additional assumption that the nonparametric function $\alpha(S)$ be continuous at each cutoff point (assumption A2*), each of the three ratios will identify a local treatment effect $\alpha(\bar{S}_j)$. Similarly, in the case where treatment effects vary randomly across individuals, it can again be shown that under assumptions A1, A2*, and the local independence assumption A3, the expression in (9) identifies a local average treatment effect $E[\alpha_i \mid S = \bar{S}_j]$ at each cutoff value $\bar{S}_j$.[19]

In case of a constant financial aid effect, we can use an analogous two-stage regression approach, where, in the first stage, we estimate the financial aid equation (5), where $E[F_i \mid S_i] = f(S_i) + \gamma_1 \cdot 1\{S_i \geq \bar{S}_1\} + \gamma_2 \cdot 1\{S_i \geq \bar{S}_2\} + \gamma_3 \cdot 1\{S_i \geq \bar{S}_3\}$. In the second stage, we use the first-stage estimate of $E[F_i \mid S_i]$ in the control function-augmented outcome equation

$$EN_i = \beta + \alpha \cdot E[F_i \mid S_i] + k(S_i) + \epsilon_i \tag{12}$$

We can estimate separate financial aid effects at each discontinuity point, $\alpha(\bar{S}_j)$ (or $E[\alpha_i \mid \bar{S}_j]$ in case of random treatment effects), by applying a two-stage procedure to estimate each local treatment effect separately, while controlling in the outcome equation for possible discontinuities at all other discontinuity points. For example, we can estimate the aid effect at $\bar{S}_1$ by including in Equation (12) the indicators $1\{S \geq \bar{S}_2\}$ and $1\{S \geq \bar{S}_3\}$

$$EN_i = \beta + \alpha(\bar{S}_1) \cdot E[F_i \mid S_i] + \sigma_1 \cdot 1\{S \geq \bar{S}_2\} + \sigma_2 \cdot 1\{S \geq \bar{S}_3\} + k(S_i) + \epsilon_i \tag{13}$$

using a first-stage predicted value of $E[F_i \mid S_i]$ and a semiparametric specification of the control function. By including the two binary indicators, under assumptions A1 and A2, $\alpha(\bar{S}_1)$ (or $E[\alpha_i \mid \bar{S}_1]$ more generally) will be identified solely by the discontinuity in the average aid amount at $\bar{S}_1$.

Having discussed alternative estimation methods for our specific RD evaluation problem, we are ready to apply them to estimate the effect of financial aid on enrollment. However, before doing so, the next section first discusses some important features of our dataset.

4. THE DATASET

Annual information on all students who apply and are admitted to College X is stored in a large computer database. This information is obtained from three

[19] In this analysis I have maintained a linear treatment effect assumption. If the treatment dose level has a nonlinear effect on the outcome variable (in which case α will vary directly with F), then continuity assumptions will generally not be sufficient for the ratio defined in (9) to identify any *specific* value of $\alpha(F)$. However, the results of Angrist and Imbens (1995) can be extended to show that in the case of nonlinear treatment effects, under assumptions A1, A2*, and A3 the ratio will in fact identify a weighted average causal response across the different levels of F observed in the data, for those with values of S equal to the cut-off value. Similarly, if we replace assumption A3 by HTV's weaker local monotonicity assumption then, in case of a fuzzy RD design, the ratio identifies a weighted average causal response across the different treatment dose levels observed in the data, for those whose treatment dose changes discontinuously at the cut-off value.

different sources: each student's original application package, the student's financial aid application forms, and a record kept by the college of the total financial aid package offered to the student and the student's subsequent enrollment decision.

The application package for college admission provides information on a wide range of student characteristics, such as age, gender, race, place of residence, citizenship, as well as information about each student's high school record. The latter information is available in the form of transcripts reporting individual course grades, the student's high school GPA, SAT scores, school and class rank as well as letters of recommendation and a statement of purpose. For those eligible for federal or state aid, the financial aid application form contains information on the income of the student's parents, as well as their expected financial contribution. The database also includes information about the aid package offered to each student by College X, consisting of the amounts of different types of financial aid that make up the total aid package of a student—loan or grant, federal aid, or College X-sponsored aid. In our empirical analysis we will focus on the effect of local, college-sponsored aid.

The information on each year's pool of admitted applicants is stored in two different files. One file includes only those applicants (referred to as FILERS) who did submit the Free Application for Federal Financial Aid (FAFSA) form to apply for federal aid jointly with a Need Analysis Document (NAD) form to apply for aid from College X. The NAD form contains personal information not provided by the student's application form for college admission, such as reported parental income and the expected parental contribution to schooling costs. The submission of these forms makes the applicant eligible to receive federal and state funded financial aid for college, which can be provided in the form of a federal or state grant and/or in the form of a loan, where the latter typically is the largest. On the basis of their academic ability, parental income, and expected financial need, an estimate of the total amounts and types of state and federal financial aid the student will receive is calculated using a set of formulas, which is then forwarded to College X. On the basis of this estimate, the college then determines the amount of discretionary aid to offer to the student. Finally, all different sources of financial aid are combined to form the total aid package offered to the applicant.

The other data file includes all those applicants who did not formally apply for federal aid (NONFILERS). This group includes individuals who do not qualify for federal aid because of high parental income as well as foreign citizens who are not permanent residents of the United States and who are therefore ineligible for any federal aid. Unlike the data file on filers, the file on nonfilers does not include information on parental income and their expected contribution. These applicants, who are not eligible for federal and state aid, are nevertheless considered for discretionary college aid. As with the filers, the nonfilers are sent a complete aid package, together with a notice of admission. Because of differences in the financial aid allocation rule and in the information that is available, and because the enrollment decision is likely to differ for both groups, in our empirical analysis both groups will be studied separately. Moreover, because all College X's aid to nonfilers takes the form of grants whereas grants constitute 88% of all college aid

TABLE 1

SAMPLE MEANS

	Variable	Filers Total	Filers Enrol	Filers Not Enrol	Filers % Enrol	Nonfilers Total	Nonfilers Enrol	Nonfilers Not Enrol	Nonfilers % Enrol
	GPA	3.40	3.34	3.43		3.26	3.19	3.28	
	SAT	1160	1133	1171		1179	1159	1182	
	F	5080	6018	4673		1012	777	1052	
	$\%F=0$	0.12	0.07	0.14		0.73	0.77	0.72	
RANK4		0.50	0.60	0.46	36.3	0.60	0.71	0.58	17.2
	GPA	3.11	3.10	3.12		3.01	2.99	3.02	
	SAT	1093	1083	1099		1123	1119	1124	
	F	4016	5027	3439		126	50	141	
	$\%F=0$	0.13	0.07	0.17		0.96	0.98	0.95	
RANK3		0.20	0.17	0.21	26.3	0.15	0.08	0.17	7.3
	GPA	3.52	3.55	3.51		3.47	3.47	3.47	
	SAT	1168	1147	1176		1198	1207	1197	
	F	4633	5845	4200		288	692	256	
	$\%F=0$	0.14	0.13	0.15		0.91	0.85	0.91	
RANK2		0.19	0.13	0.22	20.9	0.15	0.16	0.14	16.1
	GPA	3.75	3.76	3.74		3.66	3.66	3.66	
	SAT	1220	1206	1224		1263	1250	1266	
	F	5944	7587	5511		2810	2783	2816	
	$\%F=0$	0.11	0.00	0.14		0.07	0.04	0.07	
RANK1		0.11	0.10	0.12	26.6	0.10	0.05	0.11	8.0
	GPA	3.91	3.91	3.91		3.91	3.93	3.90	
	SAT	1336	1317	1343		1366	1342	1368	
	F	9085	10195	8683		4942	4494	4981	
	$\%F=0$	0.01	0.00	0.01		0.00	0.00	0.00	
	Obs	2225	674	1551	30.3	1150	168	982	14.6

NOTES: GPA represents high school grade point average and SAT is the total SAT test score. F equals total discretionary grant aid offered by college X and RANKi represents the rank an individual is assigned on the basis of the value of the ability index S, with RANK1 representing the highest rank.

offered to filers, college aid will be defined to only include grants to permit a more informative comparison of the effects of college aid on filers and nonfilers.[20]

Table 1 presents the overall enrollment rates and average financial aid offers for filers and nonfilers for the academic year 1991–92.[21] In the table, SAT represents the sum of the verbal and mathematical SAT test scores and GPA represents the student's high school grade point average, on a 4-point scale. RANKi is an indicator of the interval into which the individual's index score S falls, with RANK4 representing the interval with the lowest scores and RANK1 the interval with the highest scores. The table shows that although financial aid increases from RANK4 to RANK1, for both filers and nonfilers, the corresponding enrollment rate changes nonmonotonically with rank.

[20] As will be discussed later, the inclusion of college loans resulted in almost identical estimates.

[21] A small number of observations with missing values for the variables used in the analysis were deleted.

The table also compares those who enrolled with those who did not enroll in 1991. Filers who enrolled were offered more financial aid, on average, than those who did not enroll, suggesting a positive effect of aid on enrollment. We also find that those who enrolled have on average somewhat lower SAT scores than those who did not enroll. In contrast, there is little difference between the two groups in the average GPA. These patterns are also found when controlling for rank. Within each rank, there is little difference in average GPA, but the average SAT scores are lower for those who enrolled. It can be expected that those with higher scores, but equal rank are less likely to enroll at College X because, although they are offered similar amounts of aid from College X, they will on average receive higher financial aid offers from other colleges.

Similar, but somewhat smaller differences in average SAT and GPA scores are found for nonfilers. Different from filers, however, nonfilers who did enroll received on average lower offers of local financial aid than nonfilers who did not enroll, which suggests a negative effect of the amount of college aid on enrollment.

Figure 3 presents a scatter diagram of financial aid offers against the calculated index S for the sample of filers. Also shown in the graph is an estimated spline smooth. This estimated spline smooth $g(S)$ minimizes the sum $\sum_{i=1}^{n}(F_i - g(S_i))^2 + \lambda \int (g''(S))^2\, dS$ over the class of all twice differentiable functions over the observed domain of S. λ represents a smoothing parameter that

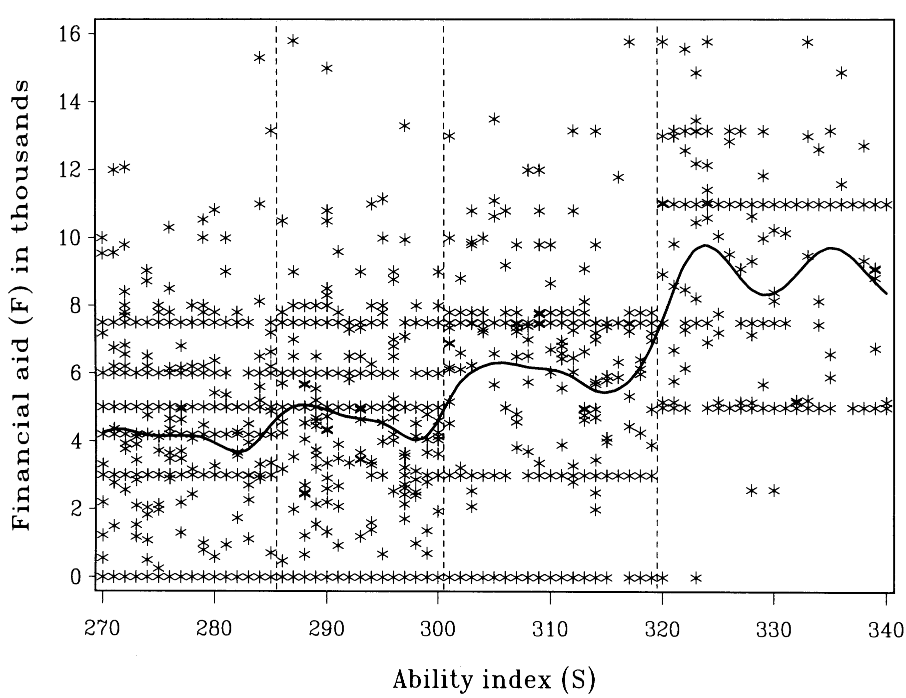

FIGURE 3

FINANCIAL AID OFFERS—FILERS. RAW DATA AND SPLINE SMOOTH (SOLID CURVE)

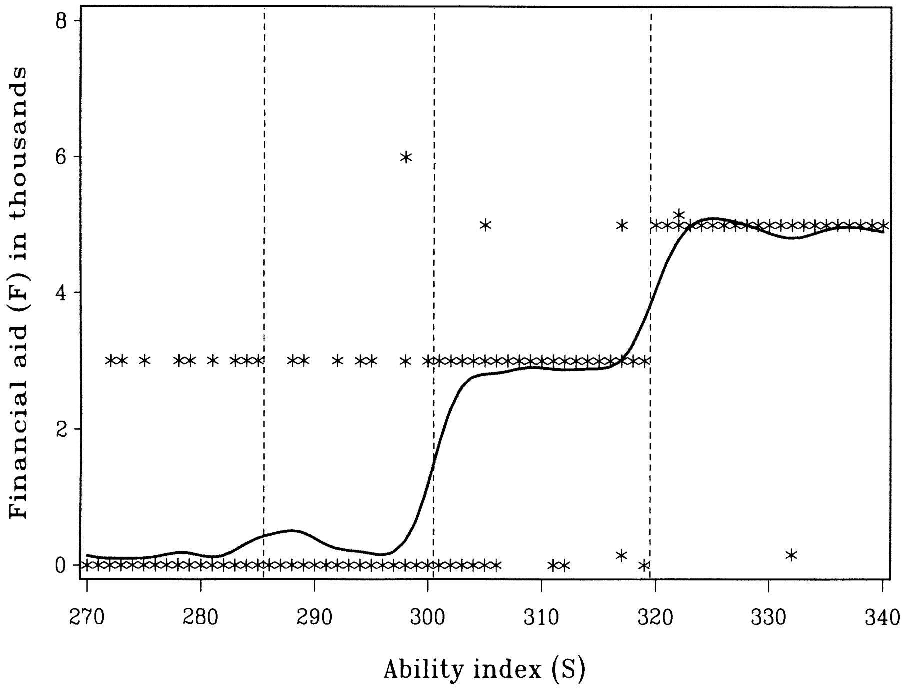

FIGURE 4

FINANCIAL AID OFFERS—NONFILERS. RAW DATA AND SPLINE SMOOTH (SOLID CURVE)

determines the weight given to the roughness penalty $\int(g''(S))^2\,dS$. The estimated curve $g(S)$ has the property that it is a cubic polynomial between two successive S-values, which, at each observation for S_i, is continuous, with continuous first and second derivatives (see Härdle, 1990, pp. 56–57). The spline smooth clearly reveals the sharp increase in the average financial aid offer at the three known cutoff points, which are represented in the figure by vertical lines. With less smoothing (lower values of the spline smoothing parameter λ), the sharpness of each increase became even more pronounced, but this also made the rest of the curve less smooth. Note that these jumps are also clearly revealed by the data themselves, with the aid values at which bunching occurs changing at each cutoff point. In the graph for nonfilers (Figure 4) these features are even more pronounced.

5. ESTIMATION RESULTS AND SENSITIVITY ANALYSIS

The financial aid allocation rule as characterized by Equation (5) and graphically displayed in Figures 3 and 4 fits the treatment allocation rule of the fuzzy RD design, with multiple cutoff points and multiple treatment (financial aid offer) levels. The average financial aid amount as a function of the ability index S contains several jumps at known cutoff values for S. Figures 5

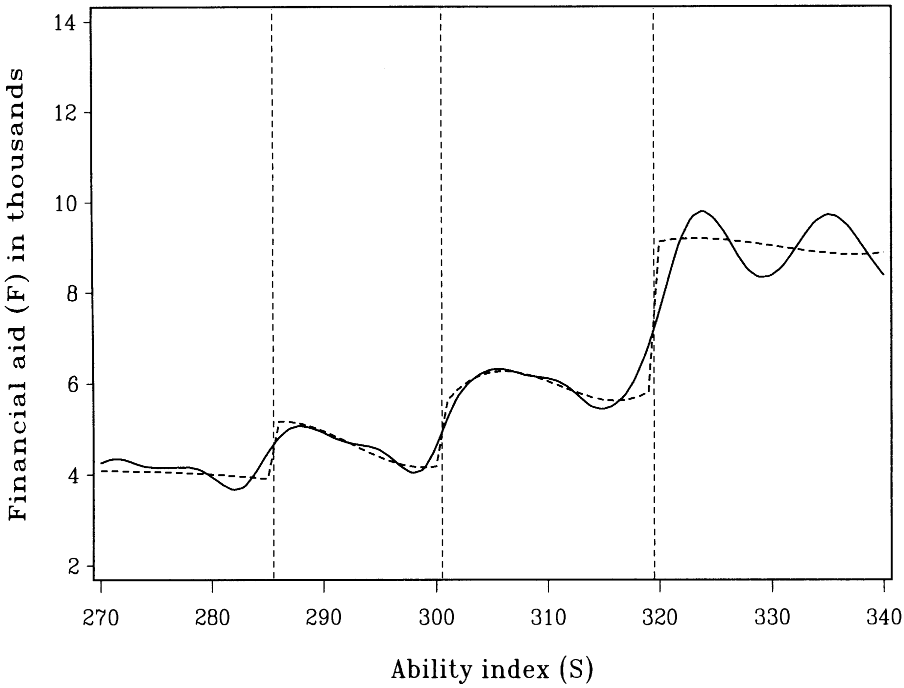

FIGURE 5

ESTIMATED FINANCIAL AID FUNCTIONS—FILERS. PIECEWISE CUBIC REGRESSION (DASHED CURVE) AND NONPARAMETRIC SPLINE SMOOTH (SOLID CURVE)

and 6 show the same spline smooth estimates presented earlier together with estimates of Equation (5) in which $f(S)$ was specified to be the continuous piecewise cubic function $f(S) = \sum_{k=0}^{3} \psi_{0k} S^k + \sum_{k=1}^{3} \psi_{1k}(S - \bar{S}_1)^k 1\{S \geq \bar{S}_1\} + \sum_{k=1}^{3} \psi_{2k}(S - \bar{S}_2)^k 1\{S \geq \bar{S}_2\} + \sum_{k=1}^{3} \psi_{3k}(S - \bar{S}_3)^k 1\{S \geq \bar{S}_3\}$. Corresponding parameter estimates of the financial aid equations are presented in Table 2.[22] Both figures, as well as the estimated discontinuities represented by the γ coefficients in Table 2, clearly reveal the importance of ability rank in the financial aid offer decision. Within each rank interval, the average aid amount offered to the group of filers declines with S, reflecting the positive and negative correlations, respectively, between the index S and family income (0.11 correlation coefficient) and between S and minority status (-0.19 coefficient) and the negative and positive effects, respectively, of income and minority status on the amount of financial aid offered. Reflecting the absence of need-based aid for nonfilers, the average amount of aid offered to nonfilers is more or less constant within rank intervals.

As a first exploration for a possible effect of financial aid offers on enrollment decisions, we can plot the enrollment rate as a function of the ability index and see whether it exhibits similar changes near the three cutoff values. Figure 7 shows

[22] The estimated regression curves for piecewise linear and quadratic specifications were almost indistinguishable from the piecewise cubic regression fits in Figures 5 and 6.

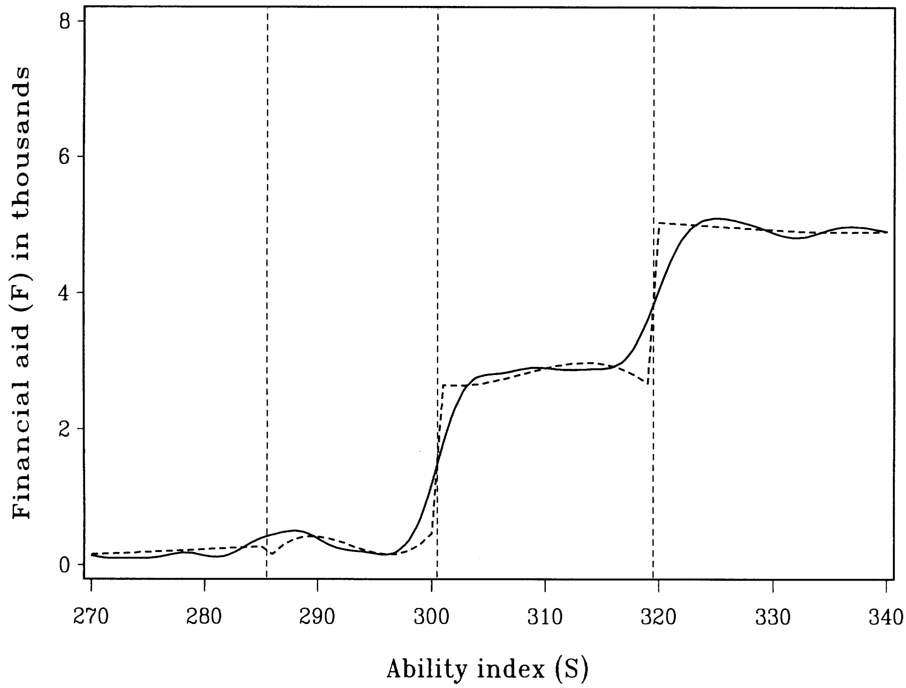

FIGURE 6

ESTIMATED FINANCIAL AID FUNCTIONS—NONFILERS. PIECEWISE CUBIC REGRESSION (DASHED CURVE) AND NONPARAMETRIC SPLINE SMOOTH (SOLID CURVE)

a spline smooth and an estimated piecewise cubic regression of the percentage enrolled as a function of the ability index score for filers. Although the slope of the spline smooth changes only slightly around the second cutoff point, it changes considerably around the other two (with a smaller value of the smoothing parameter, these changes, including that around the second cutoff point, became more pronounced). Moreover, the estimated piecewise cubic regression exhibits clear jumps at all three cutoff points. Both estimates show the enrollment rate to decline with S within a given rank interval, which is intuitive given that the average aid amount within each rank interval, as shown in Figure 5, is essentially constant (or slightly decreasing), whereas at the same time we can expect admissions and average offers from other colleges to increase with S.

Figure 8 plots the same enrollment rate functions for nonfilers. It shows a somewhat more mixed pattern of slope changes and estimated discontinuities near the cutoff points. At the first cutoff point, the enrollment rate appears to remain constant or (as suggested by the piecewise cubic regression) may actually fall slightly. However, this corresponds reasonably well with the absence of a clear jump at the first cutoff point in the average financial aid amount shown in Figure 6 and the negative insignificant estimate of γ_1 reported in Table 2. Like the average aid amount, the enrollment rate jumps up at the second cutoff point, but shows little

TABLE 2

ESTIMATES FINANCIAL AID EQUATIONS

Parameter	Filers	Nonfilers	Parameter	Filers	Nonfilers
ψ_{00}	11800	2700	ψ_{22}	−62.68	−30.51
	(14298)	(50101)		(56.06)	(31.26)
ψ_{01}	−1410	−26.31	ψ_{23}	0.267	−2.102
	(1701)	(614.6)		(2.248)	(1.230)
ψ_{02}	5.764	0.065	ψ_{31}	−156.8	146.9
	(6.873)	(2.503)		(375.7)	(175.5)
ψ_{03}	−0.008	−0.000	ψ_{32}	−41.62	16.14
	(0.009)	(0.003)		(39.80)	(18.98)
ψ_{11}	31.69	166.8	ψ_{33}	−0.989	0.488
	(254.6)	(116.5)		(1.229)	(0.605)
ψ_{12}	−18.33	−34.05	γ_1	1280	−125.3
	(42.16)	(21.96)		(487.8)	(185.7)
ψ_{13}	0.990	1.644	γ_2	1392	1959
	(1.959)	(1.073)		(752.1)	(473.6)
ψ_{21}	191.2	−287.0	γ_3	3145	2501
	(398.7)	(226.4)		(982.9)	(468.3)
Obs	2225	1150			
Adj. R^2	0.268	0.844			

NOTES: Specification: $F_i = E[F_i \mid S_i] + e_i = f(S_i) + \gamma_1 \cdot 1\{S_i \geq \bar{S}_1\} + \gamma_2 \cdot 1\{S_i \geq \bar{S}_2\} + \gamma_3 \cdot 1\{S_i \geq \bar{S}_3\} + e_i$ where $f(S)$ is the continuous piecewise cubic function $f(S) = \sum_{k=0}^{3} \psi_{0k} S^k + \sum_{k=1}^{3} \psi_{1k}(S - \bar{S}_1)^k \cdot 1\{S_i \geq \bar{S}_1\} + \sum_{k=1}^{3} \psi_{2k}(S - \bar{S}_2)^k \cdot 1\{S_i \geq \bar{S}_2\} + \sum_{k=1}^{3} \psi_{3k}(S - \bar{S}_3)^k \cdot 1\{S_i \geq \bar{S}_3\}$. Heteroskedasticity-consistent standard errors are in parentheses.

evidence of a jump around the third, where in fact the regression fit suggests a slight drop.

As discussed earlier, in case of a fuzzy RD design, under local continuity assumptions A1, A2*, and the local independence assumption A3, the ratio

$$\frac{\lim_{S \downarrow \bar{S}} E[EN \mid S] - \lim_{S \uparrow \bar{S}} E[EN \mid S]}{\lim_{S \downarrow \bar{S}} E[F \mid S] - \lim_{S \uparrow \bar{S}} E[F \mid S]} \tag{14}$$

identifies a local (average) treatment effect at each of the three cutoff points $\bar{S}_k$, $k = 1, 2, 3$. An intuitive and simple estimator for these effects is HTV's local Wald estimator. The local Wald estimator follows when one uses a one-sided uniform kernel estimator for each of the limits in (14). Selecting a band- or interval width, estimates of the financial aid effect on enrollment are obtained at each cutoff value $\bar{S}_k$ by dividing the difference in the enrollment rate between those with index scores in the interval just above the cutoff, denoted by $\overline{EN}_{\geq \bar{S}_j}$, and below the cutoff, $\overline{EN}_{<\bar{S}_j}$, with the difference in the corresponding average aid amounts $\bar{F}_{\geq \bar{S}_j} - \bar{F}_{<\bar{S}_j}$.[23]

[23] The local Wald estimator is numerically equivalent to a TSLS estimator applied to a small interval around the cutoff point $\bar{S}_j$, using the indicator $1\{S \geq \bar{S}_j\}$ as an instrument for F_i in the outcome equation (4). However, as shown in HTV, the asymptotic distributions and thus the confidence intervals of the local Wald and TSLS estimators differ. The standard errors reported in Table 3 are bootstrap standard errors.

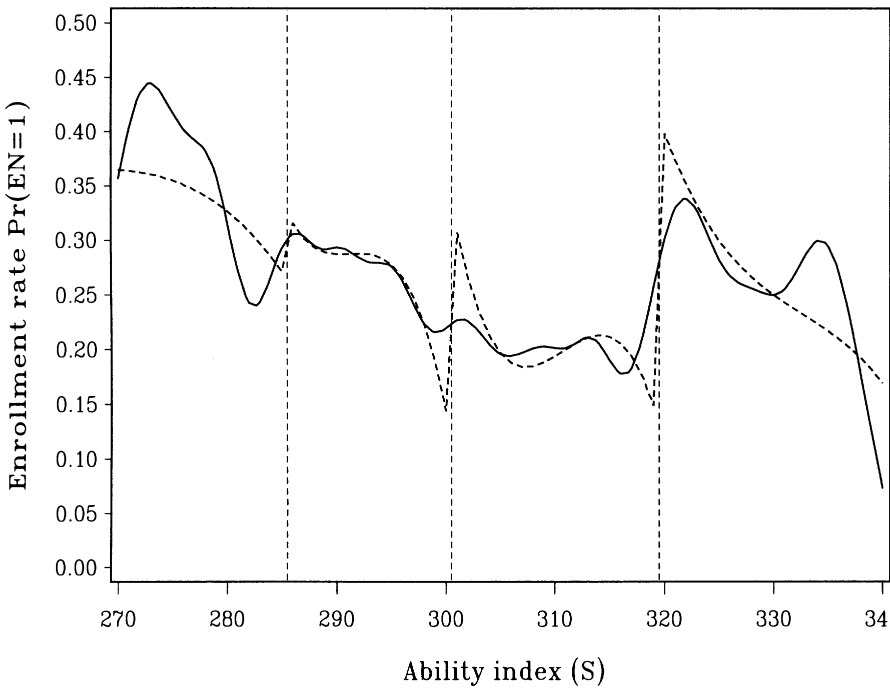

FIGURE 7

ENROLLMENT PROBABILITY—FILERS. PIECEWISE CUBIC REGRESSION (DASHED CURVE) AND NONPARAMETRIC SPLINE SMOOTH (SOLID CURVE)

As shown in Table 3, with aid measured in thousands of dollars, comparing those within three points (units of S) below and above each cutoff results in effect estimates 0.010, 0.040, and 0.067 at cutoff points $\bar{S}_1$, $\bar{S}_2$, and $\bar{S}_3$, respectively. Similarly, for nonfilers, the effect estimates are 0.523, 0.036, and −0.030. Corresponding estimates based on the much smaller samples of individuals within two points on either side of the cutoff point were 0.052, 0.075, and 0.107 for filers and 0.076, 0.060, and −0.043 for nonfilers. The relatively large standard errors for these estimates obviously reflect the modest sample sizes on which these local estimates are based. Although the effect estimate for filers seems to be increasing in S, for nonfilers the estimate appears to be decreasing in S. Note that these estimates correspond closely in direction and relative size to the financial aid effects implied by Figures 7 and 8. Also shown in Table 3 are estimates obtained by pooling the three local samples and by using the three indicators $1\{S \geq \bar{S}_j\}$, $j = 1, 2, 3$, and sample-specific intercepts as instruments in a regression of enrollment on financial aid offers and sample-specific intercepts. These pooled estimates are weighted averages of the three separate local Wald estimates, obtained by restricting all three local estimates to be the same. The pooled estimates are 0.049 and 0.088 (for three- and two-point intervals, respectively) for filers, and 0.015 and 0.033 for nonfilers.

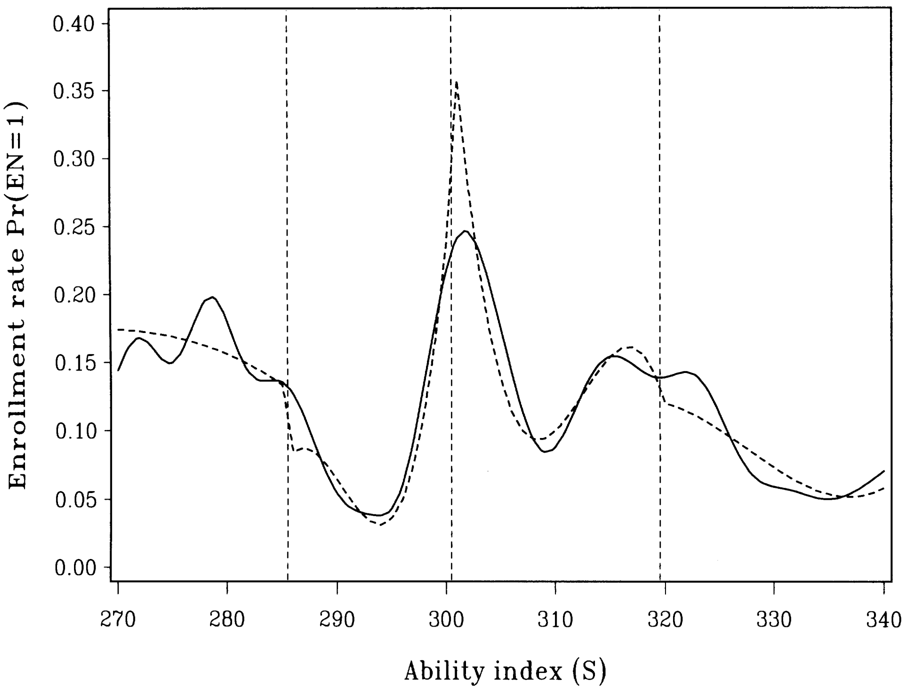

FIGURE 8

ENROLLMENT PROBABILITY—NONFILERS. PIECEWISE CUBIC REGRESSION (DASHED CURVE) AND NONPARAMETRIC SPLINE SMOOTH (SOLID CURVE)

For neither group could the hypothesis that the three local estimates are the same be rejected.

Although these nonparametric estimates are based on small samples of observations that lie within a narrow interval around a cutoff point, the semiparametric two-stage method proposed in Section 3, by relying on additional smoothness assumptions, also uses information from student applicants with ability scores outside these narrow intervals. When applied to enrollment equation (4), we estimate

$$EN_i = \beta + \alpha \cdot E[F \mid S_i] + k(S_i) + \epsilon_i \tag{15}$$

using an estimate of $E[F \mid S_i]$ from a first-stage regression and where $k(S)$ represents a power series approximation of $E[w_i \mid S]$. In the first-stage regression, the financial aid equation is specified as piecewise cubic (the most flexible specification considered here), the estimates of which were presented in Table 2.

Two-stage estimates of the enrollment equation are presented in Table 4.[24] The effect estimate for filers was found to be 0.051 whereas for nonfilers it was

[24] A Probit model specification produced estimated elasticities and overall qualitative results that were very similar to those reported here for the linear probability model. The linear probability model

TABLE 3

LOCAL WALD ESTIMATES OF FINANCIAL AID EFFECT

	Filers			Nonfilers		
	Estimate	Std Error	Obs	Estimate	Std Error	Obs
3-point intervals						
$\bar{S}_1$	0.010	(0.238)	171	0.524	(4.656)	77
$\bar{S}_2$	0.040	(0.041)	169	0.036	(4.895)	61
$\bar{S}_3$	0.067	(0.029)	107	−0.030	(0.052)	32
Pooled intervals	0.049	(0.021)	447	0.015	(0.038)	170
2-point intervals						
$\bar{S}_1$	0.052	(5.034)	109	0.076	(0.307)	53
$\bar{S}_2$	0.075	(0.049)	120	0.060	(0.052)	42
$\bar{S}_3$	0.107	(0.073)	64	−0.043	(0.045)	18
Pooled intervals	0.088	(0.028)	293	0.033	(0.040)	113

NOTES: The aid amount F is measured in thousands of 1991 dollars. Bootstrap standard errors are in parentheses. They were calculated with 10,000 bootstrap samples.

TABLE 4

TWO-STAGE ESTIMATES ENROLLMENT EQUATIONS

Variable	Filers	Nonfilers
Constant	66.19	0.722
	(44.96)	(0.192)
S	−93.36	−0.212
	(65.14)	(0.071)
S^2	49.41	
	(35.15)	
S^3	−11.55	
	(8.372)	
S^4	1.002	
	(0.743)	
F	0.051	0.019
	(0.015)	(0.011)
Observations	2225	1150

NOTES: The ability index S is measured in units of 100, and the aid amount F is measured in thousands of 1991 dollars. Heteroskedasticity-consistent standard errors are in parentheses. They have been corrected for generated regressors.

0.019. Both estimates lie in between the two pooled (local) sample-based estimates reported for filers and nonfilers in Table 3, but are more precisely estimated.[25] The effect estimate for filers corresponds to an estimated enrollment elasticity with respect to local aid evaluated at the mean of 0.86. Thus, a 10% increase in financial

has the drawback that it can generate predicted probability values outside the unit interval. However, in none of the 2225 cases for filers and 1150 cases for nonfilers did this occur.

[25] The standard errors presented in the table take into account that one of the regressors was estimated and are corrected for arbitrary forms of heteroskedasticity. However, like in most other

TABLE 5

TWO-STAGE ESTIMATES OF LOCAL FINANCIAL AID EFFECTS

	Filers		Nonfilers	
	Estimate	Std Error	Estimate	Std Error
$\bar{S}_1$	−0.010	(0.036)	−0.038	(0.149)
$\bar{S}_2$	0.017	(0.023)	0.044	(0.015)
$\bar{S}_3$	0.057	(0.016)	−0.026	(0.022)

NOTES: Heteroskedasticity-consistent standard errors are in parentheses. They have been corrected for generated regressors.

aid is predicted to lead to a 8.6% increase in the probability an average individual will enroll. The elasticity implied by the estimate for nonfilers is 0.13.[26] Given that those who did not qualify for federal financial aid are likely to be less financially constrained, it is perhaps not surprising that filers were found to be more sensitive to the offered aid amount. This result is also consistent with findings in other studies that high school graduates from higher-income families are less sensitive to financial aid offers in their college enrollment decisions compared to graduates from lower-income families (Schwartz, 1985, 1986; Leslie and Brinkman, 1988; St. John, 1990; McPherson and Schapiro, 1991). The optimal order of the power series approximation of the control function, determined by cross-validation, was found to be four, whereas for nonfilers it was one.

Although these estimates are primarily of interest to College X, the relatively large elasticity for filers indicates that financial aid is a powerful instrument for colleges to attract students. This is consistent with the dramatic increase in recent years in the use of college, and merit-based grants in particular, to compete for students. It is also consistent with recent experimental evidence from Georgia, where the introduction of the Hope scholarship apparently has had a large college enrollment effect (Dynarski, 2000).

The estimation approach used to obtain the estimates just presented assumes a constant financial aid effect. As discussed in Section 3, we can obtain separate local treatment effects by controlling in the enrollment equation for the other two discontinuities. For example, we can obtain a financial aid effect estimate at $\bar{S}_1$ by using the first-stage estimate of $E[F \mid S]$ to estimate Equation (13). Table 5 presents the corresponding RD estimates of $\alpha(\bar{S}_1)$, $\alpha(\bar{S}_2)$, and $\alpha(\bar{S}_3)$ for filers and nonfilers. The estimates are comparable to the local Wald estimates in that for filers, the estimated effect is the largest and most precisely estimated at $\bar{S}_3$, and

empirical studies using series approximation, the standard errors have not been adjusted for the fact that the order of the series approximation was determined from the data.

[26] As mentioned earlier, for comparability reasons, college aid was defined to only include local grants (which represent 92% of all local aid). When I repeated the analysis for filers using total local aid (including loans), I found almost identical results, reflecting the fact that although ability rank was important in determining grant awards, it played no role (that is, there were no discontinuities) in the determination of loan offers. Therefore, strictly speaking, the estimated effect applies to college grants only.

seems to increase with S. For nonfilers, with the exception of the effect at $\bar{S}_1$ (which is very imprecisely estimated), the estimated effects are also similar to the local Wald estimates. Although, as expected, their standard errors are smaller than those of the local Wald estimates, they remain fairly large. Note again, that in the case of random treatment effects, each estimate reported in Table 5 is an estimate of a local average treatment effect $E[\alpha_i \mid \bar{S}_j]$.

Before investigating the sensitivity of the effect estimates to alternative econometric specifications, it is important to discuss in some more detail the validity of the main assumption underlying the RD approach to evaluating the financial aid effect on enrollment: that those with values of S close to each $\bar{S}_j$ would have had similar enrollment rates if they had been offered equal amounts of aid. Formally, in terms of Equation (4), the RD approach requires local continuity of $E[w_i \mid S]$ (and $E[\alpha_i \mid S]$ in case treatment effects vary randomly across individuals) at each discontinuity point.[27] Relating this to our underlying model of enrollment decisions (Equations (1)–(3)), the first assumption requires assuming continuity of $E[u_i \mid S] = E[v_i - \delta F_i^o \mid S]$ at each of the three discontinuity points. If other colleges beside College X used exactly the same ability index S and cutoff points, it is clear that this would lead to a violation of the continuity assumption, as $E[F_i^o \mid S]$, just like $E[F_i|S]$, would be discontinuous at each cutoff point. From discussions with financial aid officers at College X and other educational institutions, it appears that the details of the financial aid rule (the definition of the ability index and rank indicators) were unique to College X. However, even if other colleges used similar cutoff schemes, as long as they did not use the exact same weighted average of the student's high school GPA and SAT test score, we can expect the average aid offers made by other colleges to be continuous in S at each cutoff point because each value of S would correspond to many different combinations of GPA, SAT, and many different income levels.[28]

Another possible violation of the local continuity assumptions could occur if students knew about the importance of the rank indicators and could manipulate their S score, for example, by retaking the SAT test if they knew their score was just below a cutoff value. If those more interested in attending College X were willing to spend additional effort to obtain a higher S, this could potentially lead to a violation of both local continuity assumptions, as the samples of individuals just below and above each cutoff score would no longer be comparable in their average characteristics, such as their unobserved preferences for College X. In that case $E[v_i \mid S]$ and therefore $E[w_i \mid S]$ in Equation (4) may become discontinuous

[27] Note that the series approximation assumes $k(S)$ to be globally continuous. This assumption can be justified by the same arguments made here for the local continuity assumption.

[28] Although several colleges and universities, such as those belonging to the Overlap group (of which College X was not a member), have in the past coordinated their financial aid policies, these agreements focused primarily on the use of a common definition of financial need and were almost exclusively concerned with need-based aid, not with merit-based grants considered here. One of the main motivations for the creation of the Overlap group was, in fact, to avoid the bidding for star students (Carlton et al., 1995).

at each cutoff and, in the case of varying treatment effects, this could also lead to a violation of assumptions A2* (continuity of the average treatment effect) and A3 (the local conditional independence assumption).[29] However, because this specific aspect of the financial aid allocation process was used solely for internal purposes and never publicized, it is unlikely that students applying to College X were aware of it. Moreover, it is unlikely that students were able to learn about the rule by analyzing offers made in previous years, as both weights ϕ_0 and ϕ_1 in the composite ability index, as well as the three cutoff points, did vary during the period studied.

Finally, without any direct control over applicants' GPA and SAT scores, and thus any linear combination of these, it is hard to imagine a reason for, or a way by which financial aid officers could generate a discontinuous change in the enrollment rates of admitted applicants with values of S just below and just above each cutoff point, which is not due to a change in financial aid. Furthermore, as in many other colleges, student admission and financial aid offer decisions in College X are made separately, and although the rank indicators play an important role in the financial aid offer decision, they were not used in the admission decision process.

It is possible, and perhaps likely, that the increase in average aid to those with a higher ability rank was targeted to individuals with specific characteristics. Although this does not lead to a violation of continuity assumptions A1 and A2*, it could lead to a violation of assumption A3 in the varying treatment effects case. HTV show, under an alternative and weaker local monotonicity assumption, that the effect estimated at a discontinuity point in this case would still represent a meaningful treatment effect: The local average treatment effect for the subgroup of individuals for whom treatment changes discontinuously at each cutoff point. For example, if the additional aid offered to those above a cutoff score was only offered to minority students, then the financial aid effect estimated here would only apply to the population of minority applicants with ability index scores near $\bar{S}$.

In summary, although in the case of heterogenous aid effects, the estimated effect may only apply to a subsample of the population studied here, we can be confident that the effect estimated here is causal and can be attributed to a change in financial aid.

5.1. *Sensitivity Analysis.* Table 6 investigates the sensitivity of our estimate of δ to alternative specifications of the financial aid equation, and of the control function in the enrollment equation. The table indicates that the estimates are relatively insensitive to the specification of the financial aid equation. This is especially true for nonfilers, which should not be surprising given that the average aid

[29] An awareness of the existence of a discontinuity in a treatment selection, eligibility, or assignment rule may compromise the validity of the RD approach in recent applications by Angrist and Lavy (1999) and Pitt and Khandker (1998). In both cases, the rules are well known and there is scope for behavioral responses of the type considered here.

TABLE 6
SENSITIVITY ANALYSIS

Specification $f(S)$ in Financial Aid Equation	Specification $k(S)$ in Enrollment Equation			
	Constant	Linear	Quadratic	Series Approximation
	Filers (2225 observations)			
Constant	−0.024	0.034	0.056	0.059
	(0.006)	(0.011)	(0.018)	(0.018)
Linear	−0.023	0.033	0.055	0.063
	(0.006)	(0.010)	(0.017)	(0.021)
Quadratic	−0.023	0.034	0.048	0.062
	(0.006)	(0.010)	(0.015)	(0.017)
Cubic	−0.023	0.033	0.046	0.062
	(0.006)	(0.010)	(0.015)	(0.017)
pw linear	−0.022	0.034	0.049	0.058
	(0.006)	(0.010)	(0.015)	(0.016)
pw quadratic	−0.022	0.033	0.046	0.055
	(0.006)	(0.010)	(0.014)	(0.016)
pw cubic	−0.023	0.032	0.043	0.051
	(0.006)	(0.010)	(0.014)	(0.015)
	Nonfilers (1150 observations)			
Constant	−0.011	0.018	0.033	0.018
	(0.006)	(0.011)	(0.019)	(0.011)
Linear	−0.011	0.019	0.036	0.019
	(0.006)	(0.011)	(0.020)	(0.011)
Quadratic	−0.011	0.019	0.035	0.019
	(0.006)	(0.011)	(0.019)	(0.011)
Cubic	−0.011	0.020	0.036	0.020
	(0.006)	(0.011)	(0.019)	(0.011)
pw linear	−0.011	0.019	0.033	0.019
	(0.006)	(0.011)	(0.019)	(0.011)
pw quadratic	−0.011	0.018	0.033	0.018
	(0.006)	(0.011)	(0.019)	(0.011)
pw cubic	-0.011	0.019	0.034	0.019
	(0.006)	(0.011)	(0.019)	(0.011)

NOTES: The financial aid equation is defined in the footnote of Table 2, and the enrollment equation is defined as $EN_i = \beta + \alpha F_i + k(S_i) + \epsilon_i$. The piecewise linear, quadratic, and cubic specifications for $f(S)$ as defined for Table 2 correspond, respectively, to the case with $K = 1$, $K = 2$, and $K = 3$. Heteroskedasticity-consistent standard errors are in parentheses. They have been corrected for generated regressors.

amount offered to nonfilers can be very accurately described by a simple piecewise constant step function as indicated by Figure 6. On the other hand, the effect estimates for both groups are quite sensitive to the specification of the control function $k(S)$. For filers, the order of the series approximation as determined by cross-validation was four, whereas for nonfilers it was one (a linear specification in S). The table indicates that an overly restrictive specification for $k(S)$ is likely to produce biased estimates.

It is interesting to compare the RD estimates in Table 4 to estimates that ignore the potential endogeneity of the financial aid amount and to estimates that

TABLE 7
COMPARISON TO OLS ESTIMATES

OLS Estimates				
No Covariates	Covariates	Quadratic in S	Covariates + Quadratic in S	RD Estimate
		Filers (2225 observations)		
0.030	0.014	0.049	0.013	0.051
(0.003)	(0.003)	(0.003)	(0.003)	(0.015)
		Nonfilers (1150 observations)		
−0.011	0.013	0.005	0.013	0.019
(0.005)	(0.010)	(0.012)	(0.014)	(0.011)

NOTES: The OLS estimates were obtained by regressing EN_i on a constant, the actual financial aid offer F_i, as well as the additional covariates listed. For filers, 15 variables were included: GPA, SAT, the individual's age, gender, U.S. citizenship, two binary indicators for race, six indicators for the applicant's state of residence, a quadratic in parental income, and a quadratic in transferable federal and state aid. For nonfilers, all variables except the parental income and federal/state aid variables were included. Heteroskedasticity-consistent standard errors are in parentheses.

try to control for its endogeneity by including a large number of observed student characteristics. Table 7 reports OLS estimates of the financial aid effect on enrollment decisions for three different specifications of the linear probability model. When relating enrollment to a constant and the actual aid amount, for filers, an OLS effect estimate of 0.030 was found, and for nonfilers, the effect estimate was −0.011. Ignoring the endogeneity of the financial aid offer in both cases leads to an underestimation of its impact on enrollment decisions. A common approach in the literature for dealing with the endogeneity problem is to include in the regression a large number of individual characteristics that are likely to be correlated with both aid offers and enrollment decisions. As indicated by the estimates reported in the second column of Table 7 for filers, controlling for observables does not necessarily lead to a reduction in the estimation bias.[30] This may reflect the remaining biases due to selection on unobservables, or may simply reflect a misspecification of the way in which the covariates enter the enrollment decision. An important attraction of the RD approach is that, by only exploiting its (local) relationship with a single observable S, one does not have to choose a functional form for the way in which other variables affect the enrollment decision.

The third and fourth columns of Table 7 present effect estimates for a specification that includes a quadratic function in S and a specification that in addition includes the set of individual characteristics. The included quadratic function of S (or for that matter a cubic or fourth-order polynomial) appears to control effectively for the omitted variables that led to the endogeneity bias, but when

[30] For filers, I included 15 variables, including GPA, SAT, the individual's age, gender, U.S. citizenship, two binary indicators for race, six indicators for the applicant's state of residence, a quadratic in parental income, and a quadratic in transferable federal and state aid. For nonfilers, all variables except the parental income and federal/state aid variables were included. Parameter estimates for these regressors are not shown for confidentiality reasons.

TABLE 8
RD EFFECT ESTIMATES, 1989–1992

	Filers		Nonfilers	
	Estimate	Obs	Estimate	Obs
1989	0.040	2182	0.014	1147
	(0.014)		(0.011)	
1990	0.041	2131	0.035	1210
	(0.019)		(0.014)	
1991	0.051	2225	0.019	1150
	(0.015)		(0.011)	
1992	0.026	2434	0.006	1169
	(0.018)		(0.010)	
1989–1992	0.041	8972	0.020	4676
	(0.007)		(0.005)	

NOTES: Entries represent the estimated coefficient on financial aid F measured in thousands of 1991 dollars. Heteroskedasticity-consistent standard errors are in parentheses. They have been corrected for generated regressors.

additional covariates are added, this is no longer the case. All together, the estimates point to the instability and unreliability of estimates that are based on the conventional approach of including the actual aid offer as an exogenous explanatory variable in an enrollment equation, providing a partial explanation for the wide range of reported aid effects in the literature.

Table 8 presents RD estimates of the financial aid effect for the years 1989 to 1992 separately and all years combined. To make the estimates comparable, all financial aid amounts were expressed in 1991 dollars. In the estimations using the pooled panel datasets, the enrollment equation (as well as the financial aid equation) included separate year-specific intercepts and slope terms for each covariate except F, the amount of financial aid offered. The effect estimates are fairly stable over time and are much larger for filers than for nonfilers in each year. The implied elasticities of enrollment with respect to aid evaluated at the pooled sample means are 0.69 for filers and 0.12 for nonfilers. If over time the alternative options of students change, this may affect their responsiveness to identical aid offers from College X in different years. If this variation in responsiveness were large, it would limit the usefulness of the effect estimate for predicting the effect of different aid policies on future enrollment. However, as shown in Table 8, I found the effect estimates to be reasonably stable during the period considered.

Finally, Tables 9 and 10 present a sensitivity analysis similar to that in Tables 6 and 7, but now for the pooled dataset. The results indicate that an overly restrictive specification of the control function is likely to produce biased estimates and, as before, the OLS-based effect estimates are found to be biased downward and to be highly sensitive to different sets of additional covariates included in the regression.

TABLE 9

SENSITIVITY ANALYSIS—EFFECT ESTIMATES, 1989–1992

Specification $f(S)$ in Financial Aid Equation	Specification $k(S)$ in Enrollment Equation: Constant	Linear	Quadratic	Series Approximation
	Filers 1989–1992 (8972 observations)			
Constant	−0.020	0.029	0.051	0.051
	(0.003)	(0.006)	(0.008)	(0.008)
Linear	−0.020	0.029	0.048	0.048
	(0.003)	(0.006)	(0.008)	(0.008)
Quadratic	−0.020	0.029	0.045	0.045
	(0.003)	(0.006)	(0.008)	(0.008)
Cubic	−0.020	0.029	0.045	0.045
	(0.003)	(0.006)	(0.008)	(0.008)
pw linear	−0.019	0.029	0.046	0.046
	(0.003)	(0.006)	(0.008)	(0.008)
pw quadratic	−0.019	0.028	0.041	0.041
	(0.003)	(0.006)	(0.007)	(0.007)
pw cubic	−0.019	0.028	0.041	0.041
	(0.003)	(0.006)	(0.007)	(0.007)
	Nonfilers 1989–1992 (4676 observations)			
Constant	−0.016	0.020	0.016	0.020
	(0.003)	(0.005)	(0.007)	(0.005)
Linear	−0.016	0.020	0.016	0.020
	(0.003)	(0.005)	(0.007)	(0.005)
Quadratic	−0.016	0.020	0.017	0.020
	(0.003)	(0.005)	(0.007)	(0.005)
Cubic	−0.016	0.020	0.017	0.020
	(0.003)	(0.005)	(0.007)	(0.005)
pw linear	−0.016	0.020	0.016	0.020
	(0.003)	(0.005)	(0.007)	(0.005)
pw quadratic	−0.016	0.020	0.016	0.020
	(0.003)	(0.005)	(0.007)	(0.005)
pw cubic	−0.016	0.020	0.016	0.020
	(0.003)	(0.005)	(0.007)	(0.005)

NOTES: See the footnote of Table 6 for definitions.

6. CONCLUSION

This article considers the important problem faced by colleges and universities of evaluating the effect of their financial aid offers on student enrollment decisions. Because of missing or incomplete information on the alternative options each student has, as well as the fact that various student characteristics influencing the subjective aid offer decision are not or cannot be kept in college databases, the college aid offer cannot be treated as an exogenous variable in an enrollment equation.

In this article, I evaluate the effect of college aid offers on student enrollment decisions at an East Coast college and show how idiosyncratic features of this college's financial aid offer process can be exploited to solve this endogeneity

TABLE 10
COMPARISON TO OLS ESTIMATES, 1989–1992

OLS estimates				
No Covariates	Covariates	Quadratic in S	Covariates + Quadratic in S	RD Estimate
Filers (8972 observations)				
0.010	0.041	0.024	0.037	0.041
(0.002)	(0.002)	(0.002)	(0.002)	(0.007)
Nonfilers (4676 observations)				
−0.012	0.019	0.016	0.018	0.020
(0.003)	(0.004)	(0.006)	(0.006)	(0.005)

NOTES: The OLS estimates were obtained by regressing EN_i on a constant, the actual financial aid offer F_i, as well as the additional covariates listed (see the footnote of Table 7 for definitions). Heteroskedasticity-consistent standard errors are in parentheses.

problem and to obtain credible effect estimates. More precisely, the college's aid assignment rule contains discontinuities that characterize the selection or assignment rule of a powerful quasi-experimental design: the RD design. In this design, the selection or assignment rule for determining who participates in a program or what treatment dose each person receives depends on a continuous variable in such a way that the corresponding propensity score, or the average treatment dose level, is discontinuous in that variable at a given cutoff point. This variable may itself be correlated with the outcome variable of interest, even in the absence of a treatment effect. In this article, I show, in the context of a traditional econometric regression model for program evaluation, how discontinuities can be exploited to identify and estimate treatment effects. A two-stage estimation procedure, is proposed and applied to estimate the effect of college aid on enrollment yields. In the first stage of this procedure, the average financial aid amount is estimated as a discontinuous function of a continuous ability index, and in the second stage, the enrollment rate is related to the estimated aid amount, while controlling for a continuous statistical relationship between the ability index and the enrollment rate.

The estimates affirm the importance of financial aid as an effective instrument in competing with other colleges for students. For applicants who had applied for federal aid (the majority of applicants), the enrollment elasticity with respect to college grants was estimated to be 0.86. The elasticity for students who were ineligible for federal aid (because of high family income or because they had no U.S. citizenship) was estimated to be much smaller, with a value of 0.13. Separate enrollment effect estimates obtained for each academic year during the 1989–1993 period indicate that this effect has been stable over time. As expected, simple OLS effect estimates were found to be biased and very sensitive to the number and types of additional covariates included in the regression to control for the endogeneity of actual aid offers.

The RD design and its ability to generate credible effect estimates exemplifies the fact that the crucial difference in avoiding bias is not whether the assignments are random or nonrandom, but whether the investigator has knowledge of, and can model, the selection process. The features of the assignment or selection rule that characterize the RD design should be relatively easy to incorporate and can, in fact, already be found in the assignment or selection rules of many other existing nonexperimental evaluation designs. In addition to the recent RD applications by Angrist and Lavy (1999) and Pitt and Khandker (1998), there are several other economic evaluation studies that have relied (at least in part) on discontinuities and/or sharp nonlinearities in the treatment assignment or selection rule in estimating treatment effects.

For example, without making an explicit connection to the design, Angrist and Krueger (1991), Imbens and van der Klaauw (1995), Yelowitz (1995), and Black (1999) all relied on the RD design in identifying and estimating causal effects. By relating the evaluation methods used in these studies to the RD approach, we are better able to evaluate the reliability of their reported estimates as well as their correct causal interpretation. As discussed in this article, the validity of the RD approach relies on several continuity assumptions, requiring that conditional on any given treatment level, the relationship between the outcome and selection variable is locally continuous. It is therefore important to assess the plausibility of this assumption in each application, especially in cases where the existence of a discontinuity in the assignment or eligibility rule is well known and individuals can alter or misreport their value of the selection variable.

The common estimation strategy in the papers listed above has essentially been to define a simple indicator for whether the selection variable is below or above some cutoff point and to use it as an instrument in TSLS or IV estimation of the outcome equation. However, because assignment in the RD design is based on a selection variable that itself may be related to the outcome variable, this indicator will generally not be a valid instrument. As discussed in this article, in the case of a RD design, treatment effects can be identified by comparing the average outcomes and treatment levels of individuals with values of the selection variable just above and below a cutoff point. By adopting on both sides of the discontinuity point a specification of the statistical relationships between the selection and treatment variable, and between the selection and outcome variable, we can use the additional information contained in observations outside this small interval. As was illustrated by the sensitivity analysis in the previous section, however, a too-restrictive specification of these relationships is likely to produce biased estimates. In general, therefore, if one suspects that the selection variable may be related to the outcome variable in the absence of a treatment effect, or if the selection rule differs from a simple piecewise-constant step function, using a common Wald estimator will be inappropriate. Similarly, applying an IV estimation procedure based on the binary instrumental variable defined earlier is likely to produce unreliable estimates, unless special care has been taken in modeling the dependence between the selection variable and treatment variable, and between the selection variable and outcome variable.

To illustrate this, consider the study by Angrist and Krueger (1991) in which quarter of birth is used as an instrument for years of schooling in an earnings equation. Quarter of birth is correlated with educational attainment because of a mechanical interaction between compulsory school attendance laws and age at school entry. The treatment assignment rule for this problem fits the fuzzy RD design, where the treatment variable is years of education and the selection variable is the individual's birthday. More precisely, the selection variable is the date the individual reaches age 6, and the state-specific date at which the student must have reached age 6 to be able to enter the first grade in that year represents the cutoff point. Those who reach age 6 just before the cutoff date are eligible to enter first grade almost a full year earlier than those who reach age 6 just after the cutoff date. Because of compulsory schooling laws requiring individuals to stay in school until they reach a given age, this discontinuity in the eligibility rule leads to a similar discontinuity in the average education level as a function of an individual's birthday at the state's cutoff date.

As most states have cutoff dates toward the end of the year, Angrist and Krueger's instrument, an indicator for whether the individual was born in the first quarter of the year or not, can be interpreted as an indicator for whether the selection variable defined above is greater or equal to the state cutoff date. The validity of the RD approach relies on an appropriate specification of the relationships between the selection and treatment variable, and between the selection and outcome variable. The selection variable in this case can be expected to be correlated with earnings even in absence of a schooling effect on earnings, given that it is correlated with the individual's age, and age is known to be correlated with earnings especially during the earlier part of the working career. Similarly, the data reveal that the average education level is not constant, but generally increases between the second and fourth quarter of each year, and one could expect similar increases to occur within each quarter.

In addition to simple Wald estimates, which ignore both dependencies, Angrist and Krueger also report TSLS estimates from specifications that include a quadratic age term (where age is measured in quarters) to control for within-year-of-birth age effects on earnings. For several of the samples, the inclusion of this quadratic age function (which can be interpreted as a specification of the control function) leads to a considerable change in the estimated returns to an additional year of schooling. Note that because age is measured in quarters, this age function is a discontinuous step function in the selection variable, which includes a discontinuity at the same value (date) at which the average education level has a discontinuity.[31] To assess the reliability of their estimates, it would therefore be useful to investigate their sensitivity to the use of quarter of birth as opposed to the continuous selection variable defined above, and to a less restrictive specification of the control function in age.[32]

[31] If the direct age effect was modeled using separate quarter of birth dummies, it would actually become impossible to estimate the schooling effect this way.

[32] The RD approach also provides an interpretation of the estimated schooling effect. As discussed earlier, in the case of varying treatment effects, under assumption A3 the estimate represents the average return to schooling for those born on the state cutoff date, whereas under the less restrictive assumption discussed in footnote 15, it identifies the effect for individuals whose educational attainment would have changed if their birthday had been just before instead of just after the state cutoff date. That is, it identifies the return to education for individuals who drop out of school as soon as they reach the compulsory education age.

Finally, the RD approach is likely to be applicable in many other policy-relevant economic evaluation studies. Many existing welfare and other government programs have benefit rules or eligibility criteria that contain discontinuities, especially programs that are means tested (such as Food Stamps) or depend on an individual's age (the now abolished AFDC program required the presence of a child below 18 years of age). Mellor and Mark (1998) discuss the potential applicability of the RD approach in studying the impact of administrative decisions involving cutoff scores, such as promotion decisions made in military and civil service organizations, as well as in salary, transfer, and layoff decisions. For example, voluntary early retirement programs (such as the buyout programs offered to the Faculty of the University of California and analyzed by Pencavel 2001) often contain pension formulas with severance incentives that contain discontinuities. The exploitation of discontinuities in these administrative programs and in the eligibility and benefit rules of various government programs to identify and estimate program effects represents a promising area for future research.

REFERENCES

ANGRIST, J. D., AND G. W. IMBENS, "Two-Stage Least Squares Estimation of Average Causal Effects in Models with Variable Treatment Intensity," *Journal of the American Statistical Association* 90 (1995), 431–42.

———, AND A. KRUEGER, "Does Compulsory School Attendance Affect Schooling and Earnings?" *Quarterly Journal of Economics* 106 (1991), 979–1014.

———, AND V. LAVY, "Using Maimonides' Rule to Estimate the Effect of Class Size on Scholastic Achievement," *Quarterly Journal of Economics* 114 (1999), 533–75.

BARNES, G. T., AND J. L. NEUFELD, "The Predictability of College Financial Aid Offers: Evidence from the Class of 1972," *Economic Inquiry* 18 (1980), 667–91.

BARNOW, B. S., G. G. CAIN, AND A. S. GOLDBERGER, "Issues in the Analysis of Selectivity Bias," in E. Stormsdorfer and G. Farkas eds., *Evaluation Studies,* Volume 5 (Beverly Hills: Sage Publications, 1980), 43–59.

BERK, R. A., AND D. RAUMA, "Capitalizing on Nonrandom Assignment to Treatments: A Regression-Discontinuity Evaluation of a Crime-Control Program," *Journal of the American Statistical Association* 78 (1983), 21–7.

BLACK, S. E., "Do Better Schools Matter? Parental Valuation of Elementary Education," *Quarterly Journal of Economics* 114 (May 1999), 577–99.

CAIN, G., "Regression and Selection Models to Improve Nonexperimental Comparisons," in C. Bennett and A. Lumsdaine eds., *Evaluation and Experiment* (New York: Academic Press, 1975), 297–317.

CAMPBELL, D. T., "Reforms as Experiments," *American Psychologist* 24 (1969), 409–29.

———, AND J. C. STANLEY, *Experimental and Quasi-Experimental Designs for Research* (Chicago: Rand-McNally, 1963).

CARLTON, D. W., G. E. BAMBERGER, AND R. J. EPSTEIN, "Antitrust and Higher Education: Was there a Conspiracy to Restrict Financial Aid?" *Rand Journal of Economics* 26 (1995), 131–47.

CARTER, G. M., J. D. WINKLER, AND A. K., BIDDLE, "An Evaluation of the NIH Research Career Development Award," Document No. R-3568-NIH, RAND, Santa Monica, CA, 1987.

COOK, T. D., AND D. T. CAMPBELL, *Quasi-Experimentation: Design and Analysis Issues for Field Settings* (Chicago: Rand-McNally, 1979).

DYNARSKI, S., "Hope for Whom? Financial Aid for the Middle Class and Its Impact on College Attendance," *National Tax Journal* 53 (2000), 629–61.

Ehrenberg, R. G., and D. R. Sherman, "Optimal Financial Aid Policies for a Selective University," *Journal of Human Resources* 19 (1984), 202–30.

Goldberger, A. S., "Selection Bias in Evaluating Treatment Effects: Some Formal Illustrations," Discussion Paper No. 123–72, IRP University of Wisconsin, 1972.

Hahn, J., P. Todd, and W. Van der Klaauw, "Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design," *Econometrica* 69 (2001), 201–09.

Härdle, W., *Applied Nonparametric Regression* (Cambridge: Cambridge University Press, 1990).

Heckman, J. J., "Varieties of Selection Bias," *American Economic Review Papers and Proceedings* 80 (1990), 313–18.

———, and V. Hotz, "Choosing Among Alternative Nonexperimental Methods for Estimating the Impact of Social Programs: The Case of Manpower Training," *Journal of the American Statistical Association* 84 (1989), 862–74.

———, and R. Robb, "Alternative Methods for Evaluating the Impact of Interventions," in J. Heckman and B. Singer eds., *Longitudinal Analysis of Labor Market Data* (Cambridge: Cambridge University Press, 1985).

———, and J. M. Snyder, "Linear Probability Models of the Demand for Attributes with an Empirical Application to Estimating the Preferences of Legislators," *Rand Journal of Economics* 28 (1997), S142–89.

———, H. Ichimura, and P. Todd, "Matching as an Econometric Evaluation Estimator: Evidence from Evaluating a Job Training Programme," *Review of Economic Studies* 64 (1997), 605–54.

Hoxby, C. M., "Benevolent Colluders? The Effects of Antitrust Action on College Financial Aid and Tuition," NBER Working Paper No. W7754, June 2000.

Imbens, G. W., and J. D. Angrist, "Identification and Estimation of Local Average Treatment Effects," *Econometrica* 62 (1994), 467–75.

———, and W. Van der Klaauw, "Evaluating the Cost of Conscription in the Netherlands," *Journal of Business Economics and Statistics* 13 (1995), 207–15.

Leslie, L. L., and P. T. Brinkman, *The Economic Value of Higher Education* (New York: American Council on Education: Macmillan, 1988).

Maddala, G. S., and L. Lee, "Recursive Models with Qualitative Endogenous Variables," *Annals of Economic and Social Measurement* 5 (1976), 525–45.

Manski, C. F., "Nonparametric Bounds on Treatment Effects," *American Economic Review Papers and Proceedings* 80 (1990), 319–23.

———, and D. A. Wise, *College Choice in America* (Cambridge: Harvard University Press, 1983).

McPherson, M. S., and M. O. Schapiro, "Does Student Aid Affect College Enrollment? New Evidence on a Persistent Controversy," *American Economic Review* 81 (1991), 309–18.

Mellor, S., and M. M. Mark, "A Quasi-Experimental Design for Studies on the Impact of Administrative Decisions: Applications and Extensions of the Regression-Discontinuity Design," *Organizational Research Methods* 1 (1998), 315–33.

Meyer, B. D., "Natural and Quasi-Experiments in Economics," *Journal of Business Economics and Statistics* 13 (1995), 151–61.

Miller, L. S., "College Admissions and Financial Aid Policies as Revealed by Institutional Practices," *Economic Inquiry* 19 (1981), 117–31.

Moore, R. L., A. H. Studenmund, and T. Slobko, "The Effect of the Financial Aid Package on the Choice of a Selective College," *Economics of Education Review* 10 (1991), 311–21.

Newey, W. K., J. L. Powell, and J. R. Walker, "Semiparametric Estimation of Selection Models: Some Empirical Results," *American Economic Review* 80 (1990), 324–28.

Parker, J., and J. Summers, "Tuition and Enrollment Yield at Selective Liberal Arts Colleges," *Economics of Education Review* 12 (1993), 311–24.

PENCAVEL, J., "The Response of Employees to Severance Incentives: The University of California's Faculty, 1991–94," *Journal of Human Resources* 36 (2001), 58–84.

PITT, M. M., AND S. R. KHANDKER, "The Impact of Group-Based Credit Programs on Poor Households in Bangladesh: Does the Gender of Participants Matter?" *Journal of Political Economy* 106 (1998), 958–96.

ROBINSON, A., AND T. D. STANLEY, "Teaching to Talent: Evaluating an Enriched and Accelerated Mathematics Program," *Journal of the Education of the Gifted* 12 (1989), 253–67.

ROSENBAUM, P. R., AND D. B. RUBIN, "The Central Role of the Propensity Score in Observational Studies for Causal Effects," *Biometrica* 70 (1983), 41–55.

SCHWARTZ, J. B., "Student Financial Aid and the College Enrollment Decision: The Effects of Public and Private Grants and Interest Subsidies," *Economics of Education Review* 4 (1985), 129–44.

———, "Wealth Neutrality in Higher Education: The Effects of Student Grants," *Economics of Education Review* 5 (1986), 107–17.

SEAVER, W. B., AND R. J. QUARTON, "Regression-Discontinuity Analysis of Dean's List Effects," *Journal of Educational Psychology* 66 (1976), 459–65.

SENECA, J. J., AND M. K. TAUSSIG, "The Effects of Tuition and Financial Aid on the Enrollment Decision at a State University," *Research in Higher Education* 26 (1987), 337–62.

SPIEGELMAN, C., "Estimating the Effect of a Large Scale Pretest Posttest Social Program," *Proceedings of the Social Statistics Section*, American Statistical Association (1979), 370–73.

ST. JOHN, E. P., "Price Response in Enrollment Decisions: An Analysis of the High School and Beyond Sophomore Cohort," *Research in Higher Education* 31 (1990), 161–76.

THISTLETHWAITE, D., AND D. CAMPBELL, "Regression-Discontinuity Analysis: An Alternative to the Ex Post Facto Experiment," *Journal of Educational Psychology* 51 (1960), 309–17.

TROCHIM, W. M., *Research Design for Program Evaluation: The Regression-Discontinuity Approach* (Beverly Hills: Sage Publications, 1984).

———, AND C. SPIEGELMAN, "The Relative Assignment Variable Approach to Selection Bias in Pretest-Posttest Group Designs," *Proceedings of the Survey Research Section*, American Statistical Association (1980), 376–80.

VENTI, S. F., "The Allocation of Discretionary Grant Aid," in C. F. Manski and D. A. Wise eds., *College Choice in America* Chap. 5 (Cambridge: Harvard University Press, 1983).

VISSER, R. A., AND J. DE LEEUW, "Maximum Likelihood Analysis for a Generalized Regression-Discontinuity Design," *Journal of Educational Statistics* 9 (1984), 45–60.

YELOWITZ, A. S., "The Medicaid Notch, Labor Supply, and Welfare Participation: Evidence from Eligibility Expansions," *The Quarterly Journal of Economics* 110 (1995), 909–39.