

EVALUATING THE CAUSAL EFFECT OF UNIVERSITY GRANTS ON STUDENT DROPOUT: EVIDENCE FROM A REGRESSION DISCONTINUITY DESIGN USING PRINCIPAL STRATIFICATION

BY FAN LI^{*,1}, ALESSANDRA MATTEI^{†,2} AND FABRIZIA MEALLI[†]

Duke University and University of Florence[†]*

Regression discontinuity (RD) designs are often interpreted as locally randomized experiments for units with a realized value of a pretreatment variable falling around a threshold. Motivated by the evaluation of Italian university grants, we consider a fuzzy RD design where the treatment status is based on both eligibility criteria and a voluntary application status. Resting on the fact that grant application and grant receipt statuses are post-assignment (post-eligibility) intermediate variables, we use the principal stratification framework to define causal estimands within the Rubin Causal Model. We propose a probabilistic formulation of the assignment mechanism underlying RD designs, by reformulating the Stable Unit Treatment Value Assumption (SUTVA) and making an explicit local overlap assumption for a subpopulation around the threshold. We invoke a local randomization assumption instead of the more standard continuity assumptions. We also develop a Bayesian approach to select the target subpopulation(s) with adjustment for multiple comparisons, and to draw inference for the target causal estimands within this framework. Applying the method to the data from two Italian universities, we find evidence that university grants are effective in preventing students from low-income families from dropping out of higher education.

1. Introduction. Amid the recent economic crisis in Europe, there has been a heated debate on how to arrange college students financial support, especially in terms of the instruments used, for example, loans, grants, tuition waiver. Accurate evaluation of the effectiveness of the existing financial aid systems is crucial for providing information to policy makers to choose between different instruments. In Italy state universities offer financial aid every year to a limited number of eligible freshmen. This intervention aims to give equal opportunity to achieve higher education to motivated students irrespective of their economic background. Dropout from university is a relevant phenomenon in Italy: indeed, the low rate of university graduates among Italian youths is mainly due to the high dropout rate (about 30%) rather than to a low enrollment rate. In this paper, we will investigate

Received December 2014; revised September 2015.

¹Supported by NSF Grants SES-1155697 and SES-1424688.

²Part of this paper was written when Alessandra Mattei was a senior fellow of the Uncertainty Quantification program of the U.S. Statistical and Applied Mathematical Sciences Institute (SAMSI).

Key words and phrases. Bayesian, causal effects, intermediate variables, principal stratification, randomization, regression discontinuity, university grants.

the causal effects of Italian university grants on student dropout rate, using data on first-year enrollees from two universities.

In the Italian university system, only students who both meet a prefixed eligibility criteria and apply for a grant can receive the grant, consisting of tuition waiver, free meals and accommodation, and a limited stipend of around 3000 Euros. Eligibility depends on an economic measurement of the student's family income and assets falling below or above a predetermined threshold. This allocation rule motivates us to adopt the regression discontinuity (RD) design framework for evaluation. RD design—a quasi-experimental design for causal inference—was first introduced in psychology by [Thistlethwaite and Campbell \(1960\)](#) and has become increasingly popular since the late 1990s in economics and other fields. Recent surveys can be found in [Cook \(2008\)](#), [Imbens and Lemieux \(2008\)](#), [Lee and Lemieux \(2010\)](#), [van der Klaauw \(2008\)](#). There are two general setups in RD designs, the sharp and the fuzzy RD designs. In the sharp design, the original form of the design, the treatment status is assumed to be a deterministic step function of a so-called *forcing variable* or *running variable*. All units with a realized value of the forcing variable on one side of a prefixed threshold are assigned to one regime and all units on the other side are assigned to the other regime. The basic idea underlying a RD analysis is that one can compare units with very similar values for the forcing variable, but different levels of treatment, to draw causal inference of the treatment at the threshold. Examples of sharp RD designs can be found, among others, in [Berk and de Leuw \(1999\)](#), [Lee \(2008\)](#), [Mealli and Rampichini \(2012\)](#). In the fuzzy design, the realized value of the forcing variable does not alone determine the receipt of the treatment, although a value of the forcing variable falling above or below the threshold acts as an encouragement or incentive to participate in the treatment. In those cases, the receipt of the treatment depends also on individual choices, which may confound treatment receipt. [Hahn, Todd and Van der Klaauw \(2001\)](#) establish a connection between fuzzy RD designs and instrumental variables (IV), and show that in a fuzzy RD setting one can identify the local average treatment effect [[Imbens and Angrist \(1994\)](#)] for a subpopulation of compliers at the threshold. Examples of fuzzy RD designs can be found, among others, in [Battistin and Rettore \(2008\)](#), [Garibaldi et al. \(2012\)](#), [van der Klaauw \(2002\)](#).

The Italian university grant allocation rule defines a fuzzy RD design because not all eligible students get a grant; they must apply first, and application is voluntary. Also, ineligible students may apply, even if they will not receive any grant. Comparing to standard fuzzy RD designs where only assignment (eligibility) and receipt of the treatment (grant) are available, the additional data on the application status in this study may provide valuable information for policy makers. In this article we develop a framework for RD analysis that is embedded in the Rubin Causal Model [[RCM](#), [Rubin \(1974, 1978\)](#)]. Resting on the fact that grant application and grant receipt statuses are post-assignment (post-eligibility) variables, we adopt the principal stratification framework [[Frangakis and Rubin \(2002\)](#)]—a generalization of the IV approach to noncompliance [[Angrist, Imbens and Rubin](#)

(1996), Imbens and Rubin (1997)]—to define causal estimands and lay the basis for inference.

Causal inference in RD designs is usually based on comparisons of units with close but distinct values of the forcing variable and relies on smoothness assumptions about the relationship between outcomes and the forcing variable around the threshold, which imply randomization at the single threshold value. For example, in fuzzy RDs, estimands are usually specified as ratio of differences of regression functions at the threshold, and inference generally relies on asymptotic approximations [e.g., Imbens and Lemieux (2008)]. In real applications, large-sample approximations might be unreliable due to the small sample size, and exact inference would be preferable. RD designs have been often described as leading to locally randomized experiments around the threshold [Dinardo and Lee (2011), Lee (2008), Lee and Lemieux (2010)]. Expanding on this interpretation, a recent strand of the literature [e.g., Cattaneo, Frandsen and Titiunik (2015), Sales and Hansen (2014)] is moving toward a formal and well-structured definition of the conditions under which RD designs are equivalent to local randomized experiments.

We further develop the idea of local randomization; we aim to provide a formal definition of the hypothetical experiment underlying RD designs, formalizing the assignment mechanism, that is, the process that describes which units got assigned to which treatment. The core of our framework is to assume there exists at least one subpopulation around the threshold where a *local overlap* assumption holds. For this subpopulation we explicitly introduce a *local randomization* assumption. Though our framework is not tied to any mode of inference, we choose the Bayesian approach for reasons explained later. We also develop a Bayesian hierarchical modeling approach to adjust for multiple comparisons in selecting the target subpopulation(s). Our work adds to the limited literature on Bayesian analysis of RD [Chib and Greenberg (2014), Chib and Jacobi (2015)].

In Section 2 we introduce the basic setup and the causal estimands. In Section 3 we propose a probabilistic formulation of the assignment mechanism for general RD designs, explicitly formulating the key assumptions, and elaborate it for the particular RD design used in the Italian university grants. Selection of the subpopulations where these assumptions hold is also discussed. A Bayesian approach for inferring causal effects in RD designs is developed in Section 4. We then apply the proposed approach to evaluate causal effects of Italian university grants on student dropout in Section 5. Section 6 concludes.

2. Causal estimands.

2.1. *Basic setup.* We introduce the notation in the context of Italian university grants. Let Z be the eligibility status, which is the initial assignment and plays the role of an “instrument” or an “encouragement,” as in randomized experiments with noncompliance. Consider a sample or population of N units; each can be either eligible to receive a treatment, $z = 1$, or ineligible, $z = 0$. In the Italian grants

system, eligibility depends on the value of a combined measurement of one's assets including income and properties, adjusted for family size, denoted by S . If a student, satisfying preliminary grade criteria, has a value of S falling below a predetermined threshold, for example, $s_0 = 15,000$ euros, he/she is eligible, and he/she is not eligible otherwise. That is, the eligibility status Z_i for student i is a deterministic function of S : $Z_i = \mathbf{1}(S_i \leq s_0)$, where $\mathbf{1}(\cdot)$ is the indicator function. Using the terminology in RD designs, S is the forcing variable.

Three variables were measured after each student i is assigned eligibility Z_i : the application status, the receipt of the grant and the dropout status; and, in principle, eligibility may affect them. Thus, we can define potential outcomes for these variables: for each student i ($i = 1, \dots, N$), given eligibility status z ($z = 0, 1$), let $A_i(z)$ be an indicator for the potential grant application status (equal to 1 if student i applies for a grant and 0 otherwise), $W_i(z)$ be an indicator for the potential treatment received (equal to 1 if student i receives a grant and 0 otherwise), and $Y_i(z)$ be the potential indicator for dropout (1 if student i drops out of university, 0 otherwise). This notation, with only two potential outcomes for each post-treatment variable for each unit, reflects the acceptance of the Stable Unit Treatment Value Assumption [SUTVA, Rubin (1980)], which implies that there is no interference between units and that there are no levels of the eligibility status other than zero and one. A more explicit formulation of SUTVA will be introduced in Section 3.1.

For each student i , given the observed eligibility status Z_i , the following variables are observed: $A_i^{\text{obs}} = A_i(Z_i)$, the observed application status; $W_i^{\text{obs}} = W_i(Z_i)$, the observed treatment received; and $Y_i^{\text{obs}} = Y_i(Z_i)$, the observed dropout status. The remaining potential outcomes are missing: $A_i^{\text{mis}} = A_i(1 - Z_i)$, $W_i^{\text{mis}} = W_i(1 - Z_i)$, and $Y_i^{\text{mis}} = Y_i(1 - Z_i)$. A vector of p pretreatment variables, \mathbf{X}_i , is also observed for each unit. We use boldface upper-case letters to denote the vector of values of a variable for all units from hereon. For example, $\mathbf{Z} = (Z_1, \dots, Z_N)'$, $\mathbf{A}^{\text{obs}} = (A_1^{\text{obs}}, \dots, A_N^{\text{obs}})'$.

2.2. The role of principal stratification for causal inference in fuzzy RD designs. In the RCM, a causal effect is defined as a comparison of the potential outcomes $Y_i(1)$ and $Y_i(0)$ for a common set of units. Obviously, in our study, such comparisons between potential dropout statuses only measure the effect of the eligibility status. To draw inference about the causal effect of receiving a grant, additional structure and assumptions are required. Since both the application status and receipt of the grant are post-assignment variables, we adopt the Principal Stratification framework [Frangakis and Rubin (2002)].

For each post-assignment variable, principal stratification defines a cross-classification of subjects into groups, namely, principal strata, defined by the joint potential values of that post-assignment variable under each of the assignments being compared. In our study, based on the application status A , subjects

are classified into four principal strata, $G_i \equiv (A_i(0), A_i(1))$: compliant-applicants $G_i = (0, 1) = CA$, students who would not apply if ineligible, but would apply if eligible; always-applicants $G_i = (1, 1) = AA$, students who would apply irrespective of their eligibility status; never-applicants $G_i = (0, 0) = NA$, students who would not apply irrespective of their eligibility status; and defiant-applicants $G_i = (1, 0) = DA$, students who would not apply if eligible, but would apply if ineligible. Because principal strata are not affected by assignment, we can define *population-average* causal effects conditional on the principal strata, known as principal causal effects:

$$(1) \quad \tau_g \equiv \mathbb{E}[Y_i(1) - Y_i(0) | G_i = g],$$

for $g = AA, CA, NA, DA$. Then the average causal effect of eligibility on dropout is a weighted average of these principal causal effects:

$$\mathbb{E}[Y_i(1) - Y_i(0)] = \sum_{g=AA,CA,NA,DA} \pi_g \tau_g,$$

where π_g is the proportion of units in principal stratum g .

Never-applicants and defiant-applicants never receive a grant, so for them we always observe the outcome in the absence of the grant. By contrast, for always-applicants and compliant-applicants we can observe $Y_i(1)$ for some eligible students who receive a grant and $Y_i(0)$ for some other ineligible students who do not receive a grant. Therefore, always-applicants and compliant-applicants are the only groups where we can learn information about the effect of receiving a grant in this study, and thus the corresponding principal causal effects, τ_{AA} and τ_{CA} , are the causal estimands of primary interest.

In the standard IV approach to noncompliance [Angrist, Imbens and Rubin (1996), Imbens and Rubin (1997)] as well as the standard setting of fuzzy RD designs [e.g., Imbens and Lemieux (2008)], data on application status are not utilized, either because they are not available or because they are ignored. Instead, the analysis is based on the principal strata formed by the post-assignment variable of grant receipt status. Specifically, there are four principal strata based on the joint potential grant receipt statuses, $R_i = (W_i(0), W_i(1))$: compliers $R_i = (0, 1)$, students who would receive the grant if eligible and would not receive the grant if ineligible; always-takers $R_i = (1, 1)$, students who would receive the grant regardless of eligibility; never-takers $R_i = (0, 0)$, student would not receive the grant regardless of eligibility; and defiers $R_i = (1, 0)$, students who would not receive the grant if eligible and would receive the grant if ineligible. The focus is generally on the causal effect for compliers:

$$\tau \equiv \mathbb{E}[Y_i(1) - Y_i(0) | R_i = (0, 1)].$$

We now establish the connection between these two sets of principal strata. The Italian grant assignment rule implies that $W_i(0) = 0$ for all i , as ineligible

units have no access to a grant, and $W_i(1) = 0$ if $A_i(1) = 0$, as eligible units need to apply for a grant to receive a grant. Therefore, by design, there are no always-takers or defiers, and the remaining principal strata R 's can be expressed as unions of principal strata G 's: never-takers comprise never-applicants and defiant-applicants, and compliers comprise always-applicants and compliant-applicants. As such, τ can be rewritten as the weighted average of the causal effects for always-applicants and compliant-applicants:

$$(2) \quad \tau = \mathbb{E}[Y_i(1) - Y_i(0)|G_i \in \{AA, CA\}] = \frac{\pi_{AA}\tau_{AA} + \pi_{CA}\tau_{CA}}{\pi_{AA} + \pi_{CA}}.$$

This illustrates that principal strata defined by the application status leads to a finer partition of the units than principal strata defined by the grant-receipt status. Indeed, the standard IV causal estimand—the causal effect for compliers τ —provides information on a “marginal” (with respect to application behavior) causal effect. If causal effects are homogeneous, marginalizing over application behavior does not critically affect the evaluation analysis. Conversely, if causal effects are heterogeneous, ignoring application behavior may represent a loss of potentially useful information. For example, if the grants are found to have a higher positive effect on always-applicants than on compliant-applicants, then it would be useful and cost-effective to study the characteristics of ineligible applicants and include those into the eligibility rule to allocate additional resources.

The estimands τ_{AA} , τ_{CA} and τ represent effects of eligibility, rather than effects of the receipt of a grant. However, “the receipt of a grant” is completely confounded with “the eligibility status”: $W(z) = z \times A(z) = z$ for always-applicants and compliant-applicants. To *attribute* these effects to “the receipt of a grant,” below we can make an exclusion restriction assumption:

ASSUMPTION 1 (Exclusion restriction for compliant-applicants and always-applicants). For all units with $G_i \in \{AA, CA\}$ or, equivalently, $R_i = (0, 1)$, the effect of eligibility is only through the receipt of the grant.³

³This assumption could be formalized by introducing potential outcomes of the form $Y_i(z, a, w)$, that is, potential outcomes for Y if eligibility status Z was set to z , application status A was set to a , and grant status W was set to w ; $z, a, w \in \{0, 1\}$. Specifically, in our setting, the (stochastic) exclusion restriction assumption for compliant-applicants and always-applicants would require that for each $a', a'' \in \{0, 1\}$, $\Pr(Y_i(0, a', w)|G_i \in \{AA, CA\}) = \Pr(Y_i(1, a'', w)|G_i \in \{AA, CA\})$, $w = 0, 1$. This exclusion restriction assumes that for compliant-applicants and always-applicants, that is, compliers, the potential outcome that would realize if they were eligible and received a grant (did not receive) is equal to the potential outcome that would realize if they were ineligible and received (did not receive) a grant, irrespective of their application status. The potential outcomes $Y_i(z, a, w)$ are a priori counterfactuals for units who exhibit a value of the application status, A_i^{obs} , and a value of the grant status, W_i^{obs} , under treatment z not equal to a and w , respectively, because in one specific experiment, they can never be observed for such types of units. Here we prefer to avoid potential outcomes of the form $Y_i(z, a, w)$ by focusing on observable potential outcomes $Y_i(z)$.

Assumption 1 attributes the intention-to-treat (ITT) effect for compliers to the receipt of grant rather than eligibility. A more formal version of this assumption, which requires double-indexed notation, is given in Footnote 3 and also [Imbens and Rubin \(2015\)](#) (Chapter 23, Assumption 23.4). This type of exclusion restriction is routinely made, often implicitly, in randomized experiments with full compliance [[Mealli and Pacini \(2013\)](#), [Mealli and Rubin \(2002\)](#)].

In real studies, the *sample-average* counterpart of the *population-average* estimands may also be of interest:

$$(3) \quad \tau_g^S \equiv \frac{1}{N_g} \sum_{i:G_i=g} [Y_i(1) - Y_i(0)],$$

where $g = \text{AA}, \text{CA}, \{\text{AA}, \text{CA}\}$ and N_g is the number of units in stratum g . Usually the *sample-average* effects can be estimated more precisely than their *population-average* counterparts. The subtle difference between them in Bayesian inference is explained in Section 4. More details can be found, for example, in [Imbens and Rubin \(1997\)](#), [Rubin \(1978\)](#) and [Imbens \(2004\)](#). For simplicity, we do not distinguish between population-average and sample-average estimands in the methodological discussion, but will present both estimates in the application.

3. The basis for inference.

3.1. *Probabilistic treatment assignment mechanism in RD designs.* The complex selection process in the Italian university grants system implies that the mechanism governing the receipt of the grant, which depends on both institutional and individual choices, is not ignorable. Below we introduce a probabilistic assignment mechanism underlying the RD design considered here, which is also applicable to general RD settings with minor modifications.

We first define the assignment mechanism, which is a row-exchangeable function that assigns probabilities to all 2^N possible N -dimensional vectors of assignments \mathbf{Z} , as a row-exchangeable function that assigns probabilities to all possible N -dimensional vectors of realizations of the forcing variable, \mathbf{S} , above or below the threshold value, s_0 . Formally,

$$(4) \quad \begin{aligned} & \Pr(\mathbf{Z} = \mathbf{z} | \mathbf{A}(0), \mathbf{A}(1), \mathbf{W}(0), \mathbf{W}(1), \mathbf{Y}(0), \mathbf{Y}(1), \mathbf{X}) \\ & = \Pr(\mathbf{S} \in \Lambda | \mathbf{A}(0), \mathbf{A}(1), \mathbf{W}(0), \mathbf{W}(1), \mathbf{Y}(0), \mathbf{Y}(1), \mathbf{X}), \end{aligned}$$

where $\mathbf{z} \in \{0, 1\}^N$ and $\Lambda \in \{(-\infty, s_0]^N, (-\infty, s_0]^{N-1} \times (s_0, \infty), (s_0, \infty) \times (-\infty, s_0]^{N-1}, \dots, (-\infty, s_0] \times (s_0, \infty)^{N-1}, (-\infty, s_0]^{N-1} \times (s_0, \infty), (s_0, \infty)^N\}$. Since Z is a deterministic function of S , the assignment mechanism can be formulated with respect to either Z or S . Here we prefer S because it is the underlying random variable that describes the reasons for the missing and observed values of potential outcomes: a value of S is assigned, which in turn determines a value for Z .

Statistical inference for causal effects requires assumptions on the assignment mechanism. We introduce assumptions that allow us to describe RD settings as classical randomized experiments around the threshold. The assignment mechanism in equation (4) is a classical randomized experiment if (i) it is individualistic,

$$\begin{aligned} & \Pr(\mathbf{S} \in \Lambda | \mathbf{A}(0), \mathbf{A}(1), \mathbf{W}(0), \mathbf{W}(1), \mathbf{Y}(0), \mathbf{Y}(1), \mathbf{X}) \\ &= \prod_{i=1}^n \Pr(S_i \leq s_0 | A_i(0), A_i(1), W_i(0), W_i(1), Y_i(0), Y_i(1), \mathbf{X}_i); \end{aligned}$$

(ii) it is probabilistic, which implies that for each unit, i , both events $S_i \leq s_0$ and $S_i > s_0$ have a priori a nonzero probability of occurring; (iii) it is unconfounded, that is, free of dependence of any potential outcomes; and (iv) it is a known function of its arguments.

The particular assignment rules underlying RD designs suggest that these assumptions are more reasonable for subpopulations of units who have a relatively large probability that the realized values of the forcing variable fall in a neighborhood around the threshold, s_0 . For these subpopulations, we can reasonably assume that the distribution of the forcing variable is unrelated to observed and unobserved characteristics of students. On the other hand, students with a very small (close to zero) or a very large (close to one) probability that $S_i \leq s_0$ are likely systematically different from other students. For example, potential outcomes observed for very rich students, who do not receive any grant, are plausibly different from potential outcomes for poor students with a realized value of S around the threshold, who do not receive a grant, and vice versa. Therefore, we focus on subpopulations of students who have a probability that $S_i \leq s_0$ strictly between zero and one, and sufficiently far away from zero and one. The following assumption guarantees that at least one such subpopulation of units exists.

ASSUMPTION 2 (Local overlap). Let \mathcal{U} be the random sample (or population) of units in the study. There exists a subset of units, \mathcal{U}_{s_0} , such that for each $i \in \mathcal{U}_{s_0}$, $\Pr(S_i \leq s_0) > \varepsilon$ and $\Pr(S_i > s_0) > \varepsilon$ for some sufficiently large $\varepsilon > 0$.

Assumption 2 assumes that there exists a subpopulation of units, each of whom has a nonzero probability of being assigned to either treatment level. This represents a main distinction between our framework and the existing RD literature that often describes RD designs as settings where the overlap assumption is violated. Now within the subpopulation \mathcal{U}_{s_0} we can formally introduce a modified SUTVA specific to the RD settings:

ASSUMPTION 3 (Local RD-SUTVA). For each $i \in \mathcal{U}_{s_0}$, consider two eligibility statuses $Z'_i = \mathbf{1}(S'_i \leq s_0)$ and $Z''_i = \mathbf{1}(S''_i \leq s_0)$, with possibly $S'_i \neq S''_i$. If $Z'_i = Z''_i$, that is, if either $S'_i \leq s_0$ and $S''_i \leq s_0$, or $S'_i > s_0$ and $S''_i > s_0$, then $A_i(\mathbf{Z}') = A_i(\mathbf{Z}'')$, $W_i(\mathbf{Z}') = W_i(\mathbf{Z}'')$, and $Y_i(\mathbf{Z}') = Y_i(\mathbf{Z}'')$.

Local RD-SUTVA rules out interference between units, implying that potential outcomes for a student cannot be affected by the eligibility status of other students. Local RD-SUTVA also assumes that there are no levels of the eligibility status other than zero and one. This component of RD-SUTVA implies that values of the forcing variable leading to the same eligibility status cannot alter potential outcomes for any unit, and thus allows us to avoid defining potential outcomes as functions of the forcing variable. Under the local RD-SUTVA for each unit within \mathcal{U}_{s_0} there exist only two potential outcomes for each post-assignment variable, corresponding to the realized value of the forcing variable falling *below* and *above* the threshold, respectively.

Finally, we formalize the concept of RD design as a local randomized experiment:

ASSUMPTION 4. (Local randomization). For each $i \in \mathcal{U}_{s_0}$,

$$\Pr(S_i | A_i(0), A_i(1), W_i(0), W_i(1), Y_i(0), Y_i(1), \mathbf{X}_i) = \Pr(S_i).$$

Assumption 4 states that within the subpopulation \mathcal{U}_{s_0} a Bernoulli trial has been conducted, with individual assignment probabilities depending only on the distribution of the forcing variable, $\Pr(Z_i = 1) = \Pr(S_i \leq s_0)$, but not on either the potential outcomes or pre-treatment variables. This assumption is crucial in justifying the key idea underlying any RD design. It implies that the eligibility statuses are randomly assigned in some small neighborhood, \mathcal{U}_{s_0} , around s_0 .

Assumption 4 may not always be plausible. For instance, when the forcing variable is a deterministic variable, which conceptually cannot be interpreted as a random variable with a nondegenerate probability distribution (such as time), the underlying design cannot, in general, be interpreted as a locally randomized experiment [see Section 6.3 in Lee and Lemieux (2010), page 347].

There are subtle but substantive differences between local RD-SUTVA and local randomization. Local RD-SUTVA is an exclusion restriction assumption and it is required to make the representation of potential outcomes as functions of the eligibility status adequate. Local randomization is an independence assumption that is crucial for inference. RD-SUTVA is different from independence assumptions: it does not imply that the probability that we observe a value of the forcing variable above or below the threshold does not depend on potential outcomes. RD-SUTVA simply implies that the exposure to assignment level z specifies a well-defined potential outcome, for all unit i and assignment levels z . In other words, considering potential outcomes as random variables, RD-SUTVA does not imply that potential outcomes have the same distribution for each value of the forcing variable. In order to make the forcing variable independent of potential outcomes, we need an additional assumption, such as Assumption 4.

Following Assumption 2, we can define a local version of the target estimands,

$$(5) \quad \tau_{g,s_0} \equiv \mathbb{E}[Y_i(1) - Y_i(0) | G_i = g, i \in \mathcal{U}_{s_0}],$$

for $g = \text{AA}, \text{CA}, \{\text{AA}, \text{CA}\}$ and their *finite-sample* counterparts, and we have

$$\tau_{\{\text{AA}, \text{CA}\}, s_0} \equiv \tau_{s_0} = \frac{\tau_{\text{AA}, s_0} \pi_{\text{AA}, s_0} + \tau_{\text{CA}, s_0} \pi_{\text{CA}, s_0}}{\pi_{\text{AA}, s_0} + \pi_{\text{CA}, s_0}},$$

where $\pi_{g, s_0} = \Pr(G_i = g | i \in \mathcal{U}_{s_0})$ for $g = \text{AA}, \text{CA}, \text{NA}, \text{DA}$, denote the proportion of principal strata in \mathcal{U}_{s_0} . A special case of \mathcal{U}_{s_0} contains the subpopulation of units with a realized value of the forcing variable *exactly* equal to s_0 .

Also, Assumption 4 implies that

$$\mathbb{E}[Y_i(1) - Y_i(0) | G_i = g, i \in \mathcal{U}_{s_0}] = \mathbb{E}[Y_i(1) - Y_i(0) | Z_i = 1, G_i = g, i \in \mathcal{U}_{s_0}].$$

Under the allocation rule of the Italian university grants, $Z_i = W_i^{\text{obs}}$ for always-applicants and compliant-applicants. Therefore, the local randomization assumption allows the estimands τ_{AA, s_0} , τ_{CA, s_0} and τ_{s_0} to be interpreted as causal effects of receiving a grant for subpopulations of students who actually receive a grant, analogous to the notion of average treatment effect for the treated (ATT).

3.2. Two additional assumptions. The following two assumptions—though not necessary for Bayesian inference—are plausible in our study and can sharpen the inference.

ASSUMPTION 5. Monotonicity of application status:

$$A_i(1) \geq A_i(0), \quad \text{for all } i \in \mathcal{U}_{s_0}.$$

ASSUMPTION 6. Stochastic exclusion restriction for never-applicants:

$$\Pr(Y_i(1) | G_i = \text{NA}, i \in \mathcal{U}_{s_0}) = \Pr(Y_i(0) | G_i = \text{NA}, i \in \mathcal{U}_{s_0}).$$

Monotonicity rules out the existence of defiant-applicants. The exclusion restriction rules out direct effects of eligibility on dropout for never-applicants. Never-applicants are students who would never apply for a grant irrespective of their eligibility status. These students would not receive the grant in any case. The exclusion restriction for never-applicants (Assumption 6) is of a very different nature from the exclusion restriction for compliant-applicants and always-applicants (Assumption 1): Assumption 6 has implications for inference but not for interpretation, whereas Assumption 1 is made solely for interpreting causal effects of assignment on the outcome as causal effects of treatment on the outcome. More discussions on the difference can be found in [Imbens and Rubin \(2015\)](#), Chapter 23 and [Mealli and Pacini \(2013\)](#).

3.3. *Selection of the subpopulations.* An important issue in practice is how to select the subpopulation \mathcal{U}_{s_0} . There can be a diverse choice of the shape of the subpopulation. In this paper, we limit our choice to symmetric intervals about s_0 for convenience. Specifically, we assume the following:

ASSUMPTION 7. There exists $h > 0$ such that for each $\varepsilon > 0$, $\Pr(s_0 - h \leq S_i \leq s_0 + h) > 1 - \varepsilon$, for each $i \in \mathcal{U}_{s_0}$.

Assumption 7 allows us to focus on the specific subsets of symmetric intervals among all neighborhoods of different shape around the threshold, s_0 . Note that Assumptions 2 and 7 do not imply that \mathcal{U}_{s_0} is unique. They only require that there exists at least one subpopulation, \mathcal{U}_{s_0} . Consequently, we are not interested in finding the largest h , but we only aim at determining plausible values for h .

Our approach for selecting bandwidth h exploits the “local” nature of Assumption 4, in the sense that it holds for a subset of units, but may not hold in general for other units. Therefore, under Assumption 4, covariates should be well balanced in the two subsamples defined by assignment in \mathcal{U}_{s_0} , and thus any test of the null hypothesis of no effect of assignment on covariates should fail to reject the null.

Assessing balance in the observed covariates raises problems of multiple comparisons, which may lead to a much higher than planned type I error if they are ignored [e.g., [Benjamini and Hochberg \(1995\)](#)]. We account for multiplicities using a Bayesian hierarchical mixed model, which provides an explicit method for borrowing information across covariates [e.g., [Berry and Berry \(2004\)](#), [Scott and Berger \(2006\)](#)]. Following [Berry and Berry \(2004\)](#), we use a mixture for the prior distribution of the eligibility parameters by assigning a point mass on equality of the means of the covariates between eligible and ineligible units. This Bayesian procedure provides a measure of the risk (posterior probability) that a chosen interval around s_0 defines a subpopulation of units that does not exactly match any true \mathcal{U}_{s_0} , including units for which our RD assumptions do not hold. More details are given in Section 5. [Cattaneo, Frandsen and Titiunik \(2015\)](#) also exploit balance tests of covariates to select \mathcal{U}_{s_0} , but their approach aims at selecting the *largest* subpopulation and does not account for multiple comparisons.

Conventional RD approaches based on local polynomial regression also involve bandwidth selection, but for a very different objective, namely, finding an optimal balance between precision and bias at the threshold for local polynomials [e.g., [Imbens and Kalyanaraman \(2012\)](#), [Lee and Lemieux \(2010\)](#), [Ludwig and Miller \(2007\)](#)], whereas our objective is to find a subpopulation where the RD assumptions are plausible and the selected subpopulation defines the target population.

3.4. *Mode of inference.* Once the subpopulation \mathcal{U}_{s_0} is chosen, and under the RD Assumptions 2–4, one can choose different modes of inference. For example, under the additional Assumptions 5 and 6, the average causal effect for compliers, τ_{s_0} , is nonparametrically identified and can be estimated using standard

moment-based methods. But the average causal effects for always-applicants and compliant-applicants, τ_{AA,s_0} and τ_{CA,s_0} , are only nonparametrically partially identified [Mealli and Pacini (2013)]. One can also use likelihood approaches to parametrically estimate causal effects [e.g., Frumento et al. (2012), Mercatanti (2013)] or use randomization-based inference [Cattaneo, Frandsen and Titiunik (2015)].

We choose Bayesian inference for the following reasons. First, causal inference in RD designs usually involves complex observational data, with multiple sources of uncertainties, including the missing potential outcomes; the Bayesian approach is particularly useful for accounting for uncertainties and for pooling information from the data in such complex settings. Second, RD analysis usually relies on a sample of units with values of the forcing variable close to a single point, the size of which may be small; Bayesian methods, not relying on asymptotic approximations, are attractive in dealing with small samples. Third, in the Bayesian paradigm, the missing potential outcomes are treated as random variables, and all inferences are based on the posterior distributions of causal estimands, which are functions of potential outcomes. Thus, inference about sample-average and population-average estimands can be drawn using the same inferential procedures. Finally, covariates can be easily incorporated in the Bayesian approach, which may reduce posterior variability of the estimates.

4. Bayesian inference. Our Bayesian analysis follows the seminal works of Rubin (1978) and Imbens and Rubin (1997). Below we give a brief outline for conducting Bayesian principal stratification analysis; the readers may refer to the existing literature for more details [e.g., Elliott, Raghunathan and Li (2010), Mattei, Li and Mealli (2013), Schwartz, Li and Mealli (2011)]. Throughout the discussion, we use $p(\cdot|\cdot)$ and $\theta_{\cdot|}$ to denote generic conditional distributions and the corresponding parameters, respectively.

Nine quantities are associated with each unit: $Y_i(0)$, $Y_i(1)$, $W_i(0)$, $W_i(1)$, $A_i(0)$, $A_i(1)$, \mathbf{X}_i , Z_i , S_i . Among these, S_i completely determines Z_i ; the principal stratum $G_i = (A_i(0), A_i(1))$ and S_i completely determine $(W_i(0), W_i(1))$. Therefore, inference for causal effects involves only $Y_i(0)$, $Y_i(1)$, $A_i(0)$, $A_i(1)$, \mathbf{X}_i , S_i , of which four are observed: S_i , \mathbf{X}_i , $A_i^{\text{obs}} = A_i(Z_i)$, $Y_i^{\text{obs}} = Y_i(Z_i)$, and two are unobserved: $A_i^{\text{mis}} = A_i(1 - Z_i)$, $Y_i^{\text{mis}} = Y_i(1 - Z_i)$.

Bayesian inference considers the observed values to be realizations of random variables and the unobserved values to be unobserved random variables. Let $p(\mathbf{Y}(0), \mathbf{Y}(1), \mathbf{A}(0), \mathbf{A}(1), \mathbf{X}, \mathbf{S}; \mathcal{U}_{s_0})$ denote the joint probability density function of these random variables for all units in \mathcal{U}_{s_0} . We assume this distribution is unit-exchangeable, that is, it is invariant under a permutation of the unit indices. Then, with essentially no loss of generality, by appealing to de Finetti's theorem [de Finetti (1937)], we can assume that there exists an unknown parameter vector θ , which is itself a random variable having a known prior distribution $p(\theta)$

such that

$$\begin{aligned}
 & p(\mathbf{Y}(0), \mathbf{Y}(1), \mathbf{A}(0), \mathbf{A}(1), \mathbf{X}, \mathbf{S}; \mathcal{U}_{s_0}) \\
 &= \int \prod_{i \in \mathcal{U}_{s_0}} p(Y_i(0), Y_i(1), A_i(0), A_i(1), \mathbf{X}_i, S_i | \boldsymbol{\theta}) p(\boldsymbol{\theta}) d\boldsymbol{\theta}.
 \end{aligned}$$

Bayesian inference of the causal estimands, which are functions of $Y_i(z)$'s and $A_i(z)$'s, centers around deriving the posterior distribution for the parameter vector of their distribution, denoted by $\boldsymbol{\theta}_{Y,G}$. Under Assumption 4, and assuming the parameters governing the distributions of the covariates, the forcing variable and the potential outcomes are a priori distinct and independent, the posterior distribution of $\boldsymbol{\theta}_{Y,G}$ can be written as follows:

$$\begin{aligned}
 & p(\boldsymbol{\theta}_{Y,G} | \mathbf{Y}^{\text{obs}}, \mathbf{A}^{\text{obs}}, \mathbf{X}, \mathbf{S}; \mathcal{U}_{s_0}) \\
 (6) \quad & \propto p(\boldsymbol{\theta}_{Y|G}) \times p(\boldsymbol{\theta}_G) \\
 & \times \prod_{i \in \mathcal{U}_{s_0}} \left[\int \int p(Y_i(0), Y_i(1) | G_i, \mathbf{X}_i; \boldsymbol{\theta}_{Y|G}) p(G_i | \mathbf{X}_i; \boldsymbol{\theta}_G) dY_i^{\text{mis}} dA_i^{\text{mis}} \right].
 \end{aligned}$$

The above decomposition suggests that two models need to be specified for model-based inference: (i) the model for potential outcomes conditional on principal strata and covariates, and (ii) the model for principal strata conditional on covariates, as well as the prior distribution for the parameters, $p(\boldsymbol{\theta}_{Y,G})$, with $\boldsymbol{\theta}_{Y,G} = (\boldsymbol{\theta}_G, \boldsymbol{\theta}_{Y|G})$.

Let $\pi_{i,g} = \Pr(G_i = g | \mathbf{X}_i; \boldsymbol{\theta}_G)$ and $f_{i,gz} = p(Y_i(z) | G_i = g, \mathbf{X}_i; \boldsymbol{\theta}_{Y|G})$. Then the posterior distribution of $\boldsymbol{\theta}_{Y,G}$ given the observed data can be written as follows:

$$\begin{aligned}
 & p(\boldsymbol{\theta}_{Y,G} | \mathbf{Y}^{\text{obs}}, \mathbf{A}^{\text{obs}}, \mathbf{X}, \mathbf{S}; \mathcal{U}_{s_0}) \\
 (7) \quad & \propto p(\boldsymbol{\theta}_{Y,G}) \times \prod_{i \in \mathcal{U}_{s_0}: S_i > s_0, A_i^{\text{obs}} = 0} (\pi_{i,CA} f_{i,CA,0} + \pi_{i,NA} f_{i,NA}) \\
 & \times \prod_{i \in \mathcal{U}_{s_0}: S_i > s_0, A_i^{\text{obs}} = 1} \pi_{i,AA} f_{i,AA,0} \times \prod_{i \in \mathcal{U}_{s_0}: S_i \leq s_0, A_i^{\text{obs}} = 0} \pi_{i,NA} f_{i,NA} \\
 & \times \prod_{i \in \mathcal{U}_{s_0}: S_i \leq s_0, A_i^{\text{obs}} = 1} (\pi_{i,AA} f_{i,AA,1} + \pi_{i,CA} f_{i,CA,1}),
 \end{aligned}$$

where $f_{i,NA} = f_{i,NA,0} = f_{i,NA,1}$ by the exclusion restriction (Assumption 6). The likelihood function, specified by the four products, does not depend on the association between the potential outcomes $Y_i(0)$ and $Y_i(1)$. Therefore, the posterior distribution of the association parameters equal their prior distribution as long as the association parameters are a priori independent of the other parameters, as we assume henceforth. The population-average causal estimands τ_{AA,s_0} , τ_{CA,s_0} and

τ_{s_0} are functions of the parameter vector $\theta_{Y,G}$, which is free of the association parameters, therefore, inference for them does not involve the association parameters [also see discussion in [Imbens and Rubin \(1997\)](#)]. Inference for sample-average causal estimands for the units in the study, on the other hand, does generally involve the association parameters. In our application inference for sample-average causal estimands is drawn under the assumption that for each unit i , potential outcomes, $Y_i(0)$ and $Y_i(1)$, are independent conditional on \mathbf{X}_i and θ .

Posterior inference of $\theta_{Y,G}$ can be obtained using Gibbs sampling with a data augmentation step to impute the missing A_i^{mis} , iteratively drawing from the two posterior predictive distributions, $p(\theta_{Y,G} | \mathbf{Y}^{\text{obs}}, \mathbf{A}^{\text{obs}}, \mathbf{A}^{\text{mis}}, \mathbf{X}, \mathbf{S}; \mathcal{U}_{s_0})$ and $p(\mathbf{A}^{\text{mis}} | \mathbf{Y}^{\text{obs}}, \mathbf{A}^{\text{obs}}, \mathbf{X}, \mathbf{S}, \theta_{Y,G}; \mathcal{U}_{s_0})$.

Specification of $\pi_{i,g}$, $f_{i,gz}$ and the corresponding prior to posterior computation depends on the specific application. Details of the models and computation in our application will be provided in Section 5. As a general guideline, we recommend to specify $\pi_{i,g}$ and $f_{i,gz}$ conditional on both covariates \mathbf{X} and the forcing variable S , even though equation (6) suggests conditioning on S is not required. Indeed, if the true subpopulations \mathcal{U}_{s_0} were known, in theory, we would not need to adjust for S , because local randomization guarantees that for units in \mathcal{U}_{s_0} values of the forcing variable falling above or below the threshold are independent of the potential outcomes. However, in practice, the true subpopulations \mathcal{U}_{s_0} are usually unknown and once a subpopulation has been selected, that is, once a value for h , say h^* , has been chosen, there may be some units with a realized value of S between $s_0 - h^*$ and $s_0 + h^*$ who do not belong to \mathcal{U}_{s_0} . For these units there may be a relationship between the forcing variable and potential outcomes, and these potential dependences need to be modeled. Specifically, systematic differences in the forcing variable S that, by definition, occur between eligible and ineligible units may affect inference in the presence of students who do not belong to \mathcal{U}_{s_0} .

5. Evaluation of Italian university grants.

5.1. *Data.* We apply the proposed method to the data from the cohort of first-year students enrolled in 2004 to 2006 at University of Pisa and University of Florence. For each student, information on grant application status (A_i^{obs}), grant receipt status (W_i^{obs}) at the beginning of the academic year, dropout status at the end of the academic year, and covariates (\mathbf{X}_i) is obtained from the ministry of education and university administrative records. The forcing variable S is a combined economic measure of each student, calculated from one's income tax return and property adjusted for family size based on a formula that is typically not fully known to the students. In all three years, the threshold of eligibility is $s_0 = 15,000$ euros. Thus, the eligibility status (Z) is also observed. Typically, students need support from fiscal experts to compute their value of S , and the income revenue authority conducts random inspections to verify that the official tax returns are reported. These factors make it extremely difficult, if not impossible, for students

TABLE 1

Summary statistics of the first-year students enrolled in 2004–2006 at Universities of Pisa and Florence, for the students with $S_i \in (14,000, 16,000)$ euros (i.e., $h = 1000, s_0 = 15,000$)

Variable	$Z = 0$		$Z = 1$	
	$A^{\text{obs}} = 0$	$A^{\text{obs}} = 1$	$A^{\text{obs}} = 0$	$A^{\text{obs}} = 1$
Sample Size	657	304	703	444
Dropout	0.36	0.50	0.35	0.36
S (euros)	15,495	15,509	14,504	14,499
Female	0.59	0.61	0.60	0.55
HS grade	80.80	84.35	80.17	84.47
University (Pisa)	0.37	0.51	0.37	0.51

to manipulate the value of S in order to end up on the right side of the threshold. Therefore, we argue that the local randomization assumption is reasonable here. Ineligible students apply either because they are not fully aware of their eligibility status or because they hope that their application will be still considered because of extra funding or other considerations.

Covariates include sex, high school grade, high school type (4 categories), major in university (6 categories), indicator of year of enrollment (2004, 2005, 2006) and indicator of university (Pisa vs. Florence). The data only include students who had a high school grade of at least 70/100 and applied either for a grant or for a reduction of tuition fee. Summary statistics of important variables for the students with S within 1000 euros of the threshold are given in Table 1. An unadjusted comparison would suggest that the applicants have higher high-school grades, which is an important indicator of a student's academic performance, but also higher dropout rate regardless of their eligibility status.

Application rate and dropout rate as a function of S among the students are given in Figure 1. The overall dropout rate is high, consistently between 30% to 50% regardless of the economic measure. From the fitted lines using local logistic polynomial models of order 3 on the two sides of the threshold, discontinuity is clearly visible in both application rate and dropout rate at the threshold. As the economic measure increases, application rate steadily decreases, while the trend in dropout rate has a concave change at the threshold, increasing on the left of the threshold and decreasing on the right.

5.2. Selection of the subpopulation. We apply the approach discussed in Section 3.3 to find subpopulations of units where our RD assumptions hold. Specifically, we use a hierarchical Bayesian model for assessing the balance of the covariates between eligibility groups. We specify probit models for binary variables, conditional probit models for categorical variables, and Gaussian models for continuous variables. Formally, we assume that $X_j \sim N(\gamma_{0j} + \gamma_{1j}Z_i, \sigma_j^2)$ if

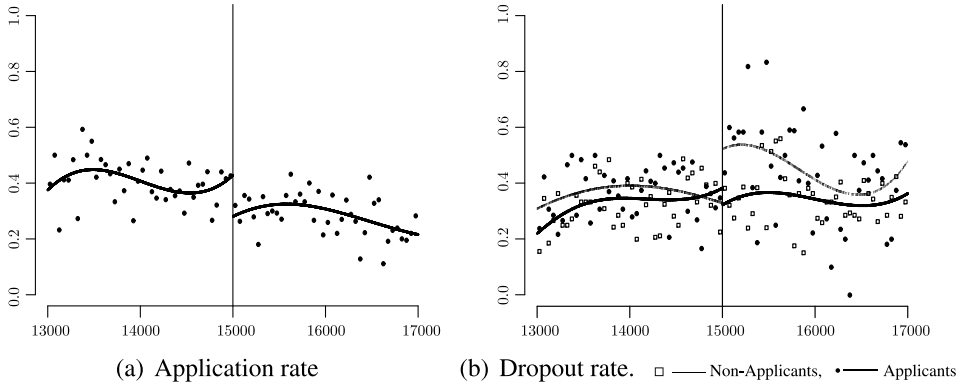


FIG. 1. Application rate (a) and dropout rate (b) as a function of the forcing variable for the first-year students in Universities of Florence and Pisa in 2004–2006. The smoothed lines are estimated using polynomial logistic regression models (of order 3) on each side of the threshold separately; each point is calculated from the units within a binwidth of 50 euros.

X_j is continuous, and $\Pr(X_{ij} = 1) = \Pr(X_{ij}^* > 0)$ with $X_{ij}^* \sim N(\gamma_{0j} + \gamma_{1j}Z_i, 1)$ if X_j is binary. If X_j is a categorical variable taking on K values, we assume that $\Pr(X_{ij} = 1) = \Pr(X_{ij}^{*(1)} \leq 0)$, and $\Pr(X_{ij} = k) = \Pr(\bigcap_{\ell=1}^{k-1} \{X_{ij}^{*(\ell)} > 0\} \cap X_{ij}^{*(k)} \leq 0)$ for $k = 2, \dots, K - 1$, where $X_{ij}^{*(k)} \sim N(\gamma_{0j}^{(k)} + \gamma_{1j}^{(k)}Z_i, 1)$, $k = 1, \dots, K - 1$, independently. Let $\boldsymbol{\gamma}_{0j} = (\gamma_{0j}^{(1)}, \dots, \gamma_{0j}^{(K-1)})'$ and $\boldsymbol{\gamma}_{1j} = (\gamma_{1j}^{(1)}, \dots, \gamma_{1j}^{(K-1)})'$.

We specify the following prior distributions for the model parameters. The variances of the continuous variables have an inverse-Gamma distribution: $\sigma_j^2 \sim \text{IG}(\underline{a}, \underline{b})$. The γ_0 's have Gaussian prior distributions: for continuous and binary variables, $\gamma_{0j} \sim N(\mu_{\gamma_0}, \sigma_{\gamma_0}^2)$, and for categorical variables, $\boldsymbol{\gamma}_{0j} \sim N(\mu_{\gamma_0} \mathbf{u}_{K-1}, \sigma_{\gamma_0}^2 \mathbf{I}_{K-1})$ with \mathbf{u}_{K-1} and \mathbf{I}_{K-1} being the $K - 1$ -dimensional vector of ones and the identity matrix of order $K - 1$, respectively. Further, for continuous and binary variables, parameters γ_{1j} are the difference between means/proportions for eligible and ineligible units. If $\gamma_{1j} = 0$, then X_j has the same distribution for eligible and ineligible units. For a categorical variable taking on K values, the proportion of units in each category is the same for eligible and ineligible units if and only if $\gamma_{1j}^{(k)} = 0$ for each $k = 1, \dots, K - 1$. We assign positive probability to these possibilities using the following mixture prior distributions:

$$\gamma_{1j} \sim \pi_{\gamma_1} \delta_0(\gamma_{1j}) + (1 - \pi_{\gamma_1}) N(\mu_{\gamma_1}, \sigma_{\gamma_1}^2)$$

and

$$\boldsymbol{\gamma}_{1j} \sim \prod_{k=1}^{K-1} [\pi_{\gamma_1} \delta_0(\gamma_{1j}^{(k)}) + (1 - \pi_{\gamma_1}) N(\mu_{\gamma_1}, \sigma_{\gamma_1}^2)],$$

where $\delta_0(\cdot)$ is the Dirac delta distribution.

For the hyperparameters, we assign the following prior distributions: $\mu_{\gamma_0} \sim N(\underline{\mu}_{\gamma_0}, \underline{\sigma}_{\gamma_0}^2)$; $\sigma_{\gamma_0}^2 \sim \text{IG}(\underline{a}_{\gamma_0}, \underline{b}_{\gamma_0})$; $\mu_{\gamma_1} \sim N(\underline{\mu}_{\gamma_1}, \underline{\sigma}_{\gamma_1}^2)$; $\sigma_{\gamma_1}^2 \sim \text{IG}(\underline{a}_{\gamma_1}, \underline{b}_{\gamma_1})$; and $\pi_{\gamma_1} \sim \text{Beta}(\underline{a}_{\pi}, \underline{b}_{\pi})$.

We implement the Bayesian model for assessing the balance of covariates on the two sides of the threshold for various subpopulations defined by different values of h . Details of the Monte Carlo Markov Chain (MCMC) for the posterior computation are relegated to the supplementary article [Li, Mattei and Mealli (2015)]. Table 2 shows the posterior probabilities that the covariates have the same distribution between eligible and ineligible students for the subpopulations defined by $h = 250, 500, 750, 1000, 1500, 2000, 2500, 3000, 4000, 5000$. These values show that the probability of the pre-assignment variables being well balanced is high for subpopulations defined by values of h strictly lower than 1500: the vast majority of these probabilities are larger than or close to 0.8. Note that the probabilities are in general lower among the covariates of “major in university,” suggesting these covariates may not be as balanced as other covariates. Nonetheless, nearly all these probabilities are still higher 0.6 with a single lowest probability being 0.565 (Tech major in university). For larger subpopulations, some covariates, such as the “indicator of university,” are clearly unbalanced.

Given that the risk that a chosen interval around the threshold defines a subpopulation that includes units not belonging to the target subpopulation, \mathcal{U}_{s_0} , is not zero, in order to account for the presence of these units, we conduct the subsequent analyses conditioning on both covariates and the realized values on the forcing variable. Also, we evaluate the robustness of our results conducting analyses using various values of h ($h = 500, 1000, 1500$).

5.3. *Parametric models.* For the units within the selected subpopulation \mathcal{U}_{s_0} , we assume parametric models for the outcome (f_{gz}) and principal strata (π_g). Alternative models, such as Student- t models [Chib and Jacobi (2015)] and Bayesian nonparametric models [Schwartz, Li and Mealli (2011)], can be considered. Note that although we are using parametric models, identification does not rely on parametric assumptions. The model for the principal strata of application consists of two conditional probit models:

$$\begin{aligned} \pi_{i,AA} &= \Pr(G_i^*(AA) \leq 0), \\ \pi_{i,NA} &= \Pr(G_i^*(AA) > 0 \text{ and } G_i^*(NA) \leq 0), \\ \pi_{i,CA} &= 1 - \pi_{i,AA} - \pi_{i,NA}, \end{aligned}$$

where

$$\begin{aligned} G_i^*(AA) &= \alpha_{AA,0} + \alpha_{AA}^{(S)} S_i^* + \mathbf{X}'_i \boldsymbol{\alpha}_{AA}^{(X)} + \varepsilon_{AA,i}, \\ G_i^*(NA) &= \alpha_{NA,0} + \alpha_{NA}^{(S)} S_i^* + \mathbf{X}'_i \boldsymbol{\alpha}_{NA}^{(X)} + \varepsilon_{NA,i}, \end{aligned}$$

with $\varepsilon_{AA,i} \sim N(0, 1)$, $\varepsilon_{NA,i} \sim N(0, 1)$ independently, and $S_i^* = (S_i - s_0)/1000$.

TABLE 2
Posterior probabilities that the covariates have the same distribution between eligible and ineligible students for various subpopulation

Variable	<i>h</i> = 250 (<i>n</i> = 528)	<i>h</i> = 500 (<i>n</i> = 1042)	<i>h</i> = 750 (<i>n</i> = 1577)	<i>h</i> = 1000 (<i>n</i> = 2108)	<i>h</i> = 1500 (<i>n</i> = 3166)	<i>h</i> = 2000 (<i>n</i> = 4197)	<i>h</i> = 2500 (<i>n</i> = 5159)	<i>h</i> = 3000 (<i>n</i> = 6113)	<i>h</i> = 4000 (<i>n</i> = 8061)	<i>h</i> = 5000 (<i>n</i> = 9846)
Sex	0.955	0.950	0.960	0.962	0.977	0.970	0.991	0.960	0.968	0.797
High school type (Baseline: Other)										
Humanity	0.951	0.952	0.949	0.955	0.979	0.970	0.965	0.986	0.953	0.962
Science	0.894	0.905	0.926	0.927	0.951	0.889	0.916	0.926	0.045	0.000
Tech	0.790	0.807	0.790	0.808	0.819	0.619	0.751	0.793	0.003	0.000
HS grade	0.955	0.958	0.972	0.978	0.971	0.981	0.987	0.990	0.984	0.986
Year (Baseline: 2004)										
2005	0.932	0.964	0.954	0.926	0.973	0.977	0.976	0.983	0.861	0.918
2006	0.883	0.918	0.914	0.909	0.959	0.934	0.952	0.970	0.807	0.884
University (Pisa)	0.950	0.916	0.971	0.983	0.686	0.097	0.225	0.300	0.082	0.000
Major in university (Baseline: Other)										
Humanity	0.946	0.899	0.689	0.797	0.798	0.932	0.958	0.990	0.964	0.946
Science	0.894	0.857	0.660	0.751	0.783	0.901	0.929	0.966	0.911	0.913
Social science	0.798	0.821	0.624	0.713	0.758	0.864	0.913	0.953	0.878	0.858
Bio-Med	0.728	0.776	0.604	0.677	0.736	0.837	0.889	0.926	0.839	0.832
Tech	0.632	0.634	0.565	0.624	0.699	0.794	0.863	0.876	0.719	0.453

Dropout, the primary outcome in our application, is binary. Therefore, we assume the following generalized linear outcome model with a probit link [Albert and Chib (1993)]:

$$\Pr(Y_i(z) = 1 | G_i = g, S_i, \mathbf{X}_i) = \Phi(\beta_{0,g,z} + \beta_{g,z}^{(S)} S_i^* + \mathbf{X}_i' \boldsymbol{\beta}_{g,z}^{(X)}).$$

We impose prior equality of the slope coefficients in the outcome regressions: $\boldsymbol{\beta}_{g,z}^{(X)} \equiv \boldsymbol{\beta}^{(X)}$ for $g = \text{AA, CA, NA}$ and $z = 0, 1$.

Define $\boldsymbol{\alpha}_g = [\alpha_{g0}, \alpha_g^{(S)}, \boldsymbol{\alpha}_g^{(X)}]'$, $g = \text{AA, NA}$, and $\boldsymbol{\beta}_{g,z} = [\beta_{0,g,z}, \beta_{g,z}^{(S)}]'$, $g = \text{AA, CA, NA}$; $z = 0, 1$. By Assumption 6, $\boldsymbol{\beta}_{\text{NA},0} = \boldsymbol{\beta}_{\text{NA},1}$. We assume that parameters are a priori independent and use multivariate normal prior distributions:

$$\boldsymbol{\alpha}_g \sim N(\underline{\boldsymbol{\mu}}_{\boldsymbol{\alpha}_g}; \underline{\boldsymbol{\sigma}}_{\boldsymbol{\alpha}_g}^2 \mathbf{I}), \quad \boldsymbol{\beta}_{g,z} \sim N(\underline{\boldsymbol{\mu}}_{\boldsymbol{\beta}_{g,z}}; \underline{\boldsymbol{\sigma}}_{\boldsymbol{\beta}_{g,z}}^2 \mathbf{I}), \quad \boldsymbol{\beta}^{(X)} \sim N(\underline{\boldsymbol{\mu}}_{\boldsymbol{\beta}}; \underline{\boldsymbol{\sigma}}_{\boldsymbol{\beta}}^2 \mathbf{I}),$$

where \mathbf{I} is the identity matrix. We specify flat priors setting the hyperparameters as follows: setting $\underline{\boldsymbol{\mu}}_{\boldsymbol{\alpha}_g}$, $\underline{\boldsymbol{\mu}}_{\boldsymbol{\beta}_{g,z}}$, $\underline{\boldsymbol{\mu}}_{\boldsymbol{\beta}}$ to be null vectors; and setting large prior variances $\underline{\boldsymbol{\sigma}}_{\boldsymbol{\alpha}_g}^2 = 10$, $\underline{\boldsymbol{\sigma}}_{\boldsymbol{\beta}_{g,z}}^2 = 10$, $\underline{\boldsymbol{\sigma}}_{\boldsymbol{\beta}}^2 = 10$ for $g = \text{AA, CA, NA}$; $z = 0, 1$.

5.4. Posterior computation. Upon obtaining the posterior draws of the parameters, we calculate three estimates for each causal estimand: population-average effect *within* \mathcal{U}_{s_0} and *at* s_0 , and sample-average effect within \mathcal{U}_{s_0} . The population-average effects within \mathcal{U}_{s_0} are calculated averaging the model-based dropout proportions over the empirical distribution of the pre-assignment variables and the forcing variable:

$$\frac{\sum_{i \in \mathcal{U}_{s_0}} \pi_{i,g} \Phi(\beta_{0,g,1} + \beta_{g,1}^{(S)} S_i^* + \mathbf{X}_i' \boldsymbol{\beta}^{(X)})}{\sum_{i \in \mathcal{U}_{s_0}} \pi_{i,g}} - \frac{\sum_{i \in \mathcal{U}_{s_0}} \pi_{i,g} \Phi(\beta_{0,g,0} + \beta_{g,0}^{(S)} S_i^* + \mathbf{X}_i' \boldsymbol{\beta}^{(X)})}{\sum_{i \in \mathcal{U}_{s_0}} \pi_{i,g}},$$

for $g = \text{AA, CA, \{AA, CA\}}$. The population-average effects at s_0 are calculated in a similar way, setting $S_i^* = 0$ (i.e., $S_i = s_0$) for each i . To obtain the sample-average estimates, we compute the posterior predictive distributions of the potential outcomes for each student i in \mathcal{U}_{s_0} .

5.5. Results. We conducted Bayesian analysis using $h = 500, 1000, 1500$. Posterior inference is based on 5000 draws from the posterior distributions simulated using single chains, which were run for 125,000 iterations. To assess convergence of iterative simulation methods, we calculated the Cramer–von-Mises statistic to test the null hypothesis that the sampled values come from a stationary distribution and visually inspected the trace-plots of the causal parameters (functions of model parameters). We also run multiple MCMC chains from different starting values from each h to evaluate the mixing of the chains using the

TABLE 3

Posterior median and 95% credible intervals of principal strata proportion and super-population and finite-sample causal effects on dropout for always-applicants (τ_{AA,s_0}), compliant-applicants (τ_{CA,s_0}) and their union (τ_{s_0}), for the subpopulation within different bandwidths h around the threshold

h	Population-average		Sample-average		Population-average at s_0	
	Median	95% CI	Median	95% CI	Median	95% CI
$h = 500$						
$\Pr(G_i = AA)$	0.323	(0.294; 0.355)	0.322	(0.309; 0.336)	0.320	(0.291; 0.352)
$\Pr(G_i = CA)$	0.060	(0.031; 0.105)	0.041	(0.021; 0.090)	0.058	(0.030; 0.094)
$\Pr(G_i = NA)$	0.616	(0.570; 0.650)	0.637	(0.590; 0.651)	0.621	(0.583; 0.654)
τ_{AA,s_0}	-0.153	(-0.313; -0.030)	-0.152	(-0.307; -0.038)	-0.154	(-0.298; -0.030)
τ_{CA,s_0}	0.045	(-0.170; 0.497)	0.074	(-0.256; 0.545)	0.039	(-0.169; 0.474)
τ_{s_0}	-0.116	(-0.253; -0.005)	-0.120	(-0.265; -0.009)	-0.120	(-0.245; -0.012)
$h = 1000$						
$\Pr(G_i = AA)$	0.336	(0.312; 0.365)	0.333	(0.318; 0.354)	0.335	(0.311; 0.363)
$\Pr(G_i = CA)$	0.043	(0.002; 0.086)	0.027	(0.002; 0.075)	0.043	(0.001; 0.075)
$\Pr(G_i = NA)$	0.623	(0.584; 0.652)	0.640	(0.599; 0.645)	0.625	(0.594; 0.656)
τ_{AA,s_0}	-0.161	(-0.273; -0.052)	-0.161	(-0.270; -0.057)	-0.154	(-0.259; -0.052)
τ_{CA,s_0}	0.028	(-0.745; 0.828)	0.031	(-0.778; 0.871)	0.010	(-0.918; 0.933)
τ_{s_0}	-0.132	(-0.242; -0.021)	-0.139	(-0.247; -0.034)	-0.128	(-0.229; -0.020)
$h = 1500$						
$\Pr(G_i = AA)$	0.332	(0.315; 0.349)	0.332	(0.326; 0.337)	0.329	(0.312; 0.346)
$\Pr(G_i = CA)$	0.042	(0.035; 0.077)	0.027	(0.020; 0.066)	0.042	(0.036; 0.062)
$\Pr(G_i = NA)$	0.625	(0.591; 0.642)	0.642	(0.605; 0.644)	0.628	(0.606; 0.646)
τ_{AA,s_0}	-0.183	(-0.286; -0.077)	-0.187	(-0.291; -0.085)	-0.153	(-0.247; -0.063)
τ_{CA,s_0}	0.010	(-0.304; 0.797)	0.011	(-0.207; 0.928)	0.000	(-0.154; 0.951)
τ_{s_0}	-0.153	(-0.256; -0.040)	-0.165	(-0.266; -0.057)	-0.130	(-0.217; -0.019)

Gelman–Rubin statistic [Gelman and Rubin (1992)]. The results provided no evidence against convergence.⁴

Table 3 shows posterior medians and 95% credible intervals for the principal strata proportions under monotonicity and for the causal parameters τ_{AA,s_0} , τ_{CA,s_0} , τ_{s_0} , for bandwidths ranging from 500 to 1500 euros. The results are robust across different bandwidths. The estimated proportions of the principal strata are very similar across different h : there are more than 61% never-applicants, more than

⁴We also conducted Bayesian analysis using alternative models with different order polynomials in S as well as models conditioning only on S (without using the pretreatment variables) and null models, conditioning on neither S nor the pretreatment covariates. Consistent with results found in Mealli and Rampichini (2012), higher order polynomials do not lead to substantial inferential benefits, and posterior distributions of the causal effects of interest did not substantially change with the alternative models, so here we only show the results based on models conditional on both S and the pretreatment covariates.

32% always-applicants and less than 6.5% compliant-applicants. The three estimates for the same causal parameter are also similar. The posterior distributions of the causal effect for always-applicants, τ_{AA,s_0} , and the union of always-applicants and compliant-applicants, τ_{s_0} , are centered on negative values, and the 95% credible intervals do not cover 0, irrespective of the choice of the bandwidth.

For instance, consider the finite-sample causal effects for the subpopulation within $h = 1000$ euros around the threshold (middle block of columns in Table 3). The estimated τ_{s_0} suggests a 13.9% [95% CI: (3.4%; 24.7%)] reduction in dropout rate for the students who receive the grants. The estimated τ_{AA,s_0} suggests an even stronger positive effect among the always-applicants: a 16.1% [95% CI: (5%; 27%)] reduction in dropout rate. In fact, τ_{s_0} , which is a weighted average of the effects for always-applicants and compliant-applicants, appears to be diluted by the somewhat surprising small and negative effect among the compliant-applicant. However, the data do not seem to contain much information on compliant-applicants (the estimated proportion of compliant-applicants is very small, less than 5%), and the effects were estimated with large uncertainties.

These results suggest that the current Italian university grants are effective in reducing dropout from universities among students from families with annual economic measure around 15,000 euros, and also reveal that the effects are heterogeneous among always-applicants and compliant-applicants. From a cost-effective perspective, our results suggest that it might be beneficial for the education administration to weaken the eligibility criteria (i.e., increase the threshold) to allow more applicants and thus always-applicants to get the grant. The combination of a low percentage of compliant-applicants and a high percentage of always-applicants suggests that most students with the economic measure being around the threshold who intend to apply for the grants would apply irrespective of their eligibility. From a policy perspective, this implies that the educational administration should better explain the rule of eligibility to potential applicants to discourage ineligible students from applying, and thus reduce unnecessary efforts from these students and the administration for processing these applications.

5.6. Posterior predictive model checking. Assessing the plausibility of model assumptions is critical in model-based approaches. Model checking here is not as crucial as in other model-based approaches, thanks to the randomization assumption, but it is still prudent to check the model fit since there are uncertainties in the selection of \mathcal{U}_{s_0} . We adopt Bayesian posterior predictive checks [Gelman, Meng and Stern (1996)] to assess goodness of fit of our models in the application. Posterior predictive checks evaluate goodness of fit of models by measuring the discrepancy between the observed data and replicated data simulated from its posterior predictive distribution. The particular procedure adopted here is similar to that in Mattei, Li and Mealli (2013), Section 6. Specifically, we consider three discrepancy measures aimed at assessing whether the model can preserve broad

TABLE 4
Bayesian p -values of signal, noise and SNR under different h for the model used in the application to Italian university grants

h	Principal strata	Signal	Noise	SNR
500	{CA}	0.095	0.630	0.094
	{AA}	0.254	0.325	0.254
	{AA, CA}	0.338	0.273	0.370
1000	{CA}	0.411	0.425	0.419
	{AA}	0.400	0.444	0.473
	{AA, CA}	0.493	0.335	0.518
1500	{CA}	0.208	0.444	0.210
	{AA}	0.372	0.400	0.261
	{AA, CA}	0.455	0.337	0.470

features of signal, noise and signal-to-noise ratio (SNR) in the dropout status distribution for compliant-applicants, always-applicants and the union of these two principal strata, and calculate posterior predictive p -values (PPPVs) to summarize discrepancies between the observed data and replicated data. Extreme (close to 0 or 1) PPPVs can be interpreted as evidence of lack of fit of the model in, at least some aspects of, the observed data. Further details of the procedure are relegated to the supplementary article [Li, Mattei and Mealli (2015)].

Table 4 shows the PPPVs for the model fit to the subpopulation with bandwidth of 500, 1000 and 1500 euros, respectively. The PPPVs suggest good model fit for all bandwidths, except for a slight under-fit for always-applicants in the subpopulation with $h = 500$, which is possibly due to the small sample size. We have also calculated the less conservative sampled posterior predictive p -values [Gosselin (2011), Johnson (2007)] and obtained similar conclusions.

6. Discussion. Motivated by the evaluation of Italian university grants, we propose a probabilistic formulation of the assignment mechanism for regression discontinuity designs and develop a full Bayesian approach to draw causal inference within the framework of principal stratification. In particular, we illustrate how to utilize information on application status to gain additional insights in program evaluation. Applying the method to the data from two Italian universities, we find that university grants reduce dropping out of higher education for students from low-income families and the effect size is especially pronounced for motivated students (always-applicants). Because the amount of grants is the same for all awardees in the data, we were not able to study the dose-response relation between the amount of grants and dropout, which is a topic of potentially important policy implications that merits further investigation.

The core of our approach is the local randomization assumption, which is intrinsically nontestable. Therefore, it may be worthwhile to conduct sensitivity analyses

to assess the robustness of the results with respect to violations of local randomization. To this end, we conduct further analyses deriving the posterior distributions of the causal effects under three additional model specifications: (i) a model where we specify the model for principal strata, $\pi_{i,g}$, and the conditional model for potential outcomes given principal strata, $f_{i,gz}$, conditioning on neither the forcing variable nor the pretreatment variables; (ii) a model where we specify $\pi_{i,g}$ and $f_{i,gz}$ conditioning only on the forcing variable, without including the pretreatment variables; and (iii) a model where we specify $\pi_{i,g}$ and $f_{i,gz}$ conditioning only on the pretreatment variables, without including the forcing variable. Under local randomization, adjusting inference for either the forcing variable or the pretreatment variables should not be required, therefore, we expect that results are similar across different model specifications. Indeed, results, shown in the supplementary article [Li, Mattei and Mealli (2015)], are robust across different models, suggesting that causal inference under the local randomization assumption is robust and credible.

A fundamental distinction between our approach and the previous local-regression based RD approaches lies in the role of the forcing variable in the analysis. Specifically, previous approaches generally view the forcing variable as a pre-assignment covariate rather than a random variable as in our approach. As a consequence, the standard overlap assumption, which requires that there are both treated and control units for all values of the covariates including the forcing variable, is violated. Violation of the overlap assumption implies that the conditional independence assumption, which trivially holds in RD settings, cannot be exploited directly. Instead some kind of extrapolation is required and, in order to avoid that estimates heavily rely on extrapolation, previous analyses focus on causal effects of the treatment for units *at* the threshold. Smoothness assumptions, for example, continuity of conditional regression functions of potential outcomes given the forcing variable, are usually assumed to draw inference on those causal effects. Local randomization and continuity are different assumptions, leading to different causal estimands: under continuity assumptions, units with a realized value of the forcing variable around the threshold are used to draw inference on causal effects for units *at* the threshold, whereas under local randomization, inference is drawn on causal effects for units *around* the threshold.

In the evaluation of Italian university grants, student's academic performance (measured by total credits taken or passing rate of exams) is also of great interest in policy. As illustrated by Mattei, Li and Mealli (2013) and Mercatanti, Li and Mealli (2015), jointly modeling two outcomes, dropout and academic performance in this case, would be worthwhile for both practical and inferential purposes.

After the first year, the Italian university grant assignment rule combines sequential and RD designs [Cellini, Ferreira and Rothstein (2010)]: grants are allocated both on the basis of students' family economic indicator and on the ground of their academic performance (exam scores above a certain threshold). Such complex assignment mechanisms pose challenges to causal inference, requiring new structures and assumptions; meanwhile, they also present great opportunities for

extending the existing framework to more general RD settings. One specific direction of our future research is to develop methods that combine Bayesian tools for RDs and dynamic treatment regimes [Murphy (2003), Zajonc (2012)] in the presence of multiple forcing variables [Imbens and Zajonc (2011)].

Acknowledgments. The authors are grateful to the Associate Editor and two reviewers for constructive comments, to Quanli Wang for computing support, and to Guido Imbens, Michael Hudgens and Sid Chib for helpful discussions.

SUPPLEMENTARY MATERIAL

Supplement to “Evaluating the causal effect of university grants on student dropout: Evidence from a regression discontinuity design using principal stratification” (DOI: [10.1214/15-AOAS881SUPP](https://doi.org/10.1214/15-AOAS881SUPP); .pdf). We describe in detail the Bayesian approach we used to select the subpopulations, the Markov Chain Monte Carlo (MCMC) methods used to simulate the posterior distributions of the parameters of the models, the posterior predictive checks, and the sensitivity analysis regarding local randomization described in Section 6.

REFERENCES

- ALBERT, J. H. and CHIB, S. (1993). Bayesian analysis of binary and polychotomous response data. *J. Amer. Statist. Assoc.* **88** 669–679. [MR1224394](#)
- ANGRIST, J. D., IMBENS, G. W. and RUBIN, D. B. (1996). Identification of causal effects using instrumental variables (with discussion). *J. Amer. Statist. Assoc.* **91** 444–472.
- BATTISTIN, E. and RETTORE, E. (2008). Ineligibles and eligible non-participants as a double comparison group in regression discontinuity designs. *J. Econometrics* **142** 715–730.
- BENJAMINI, Y. and HOCHBERG, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *J. Roy. Statist. Soc. Ser. B* **57** 289–300. [MR1325392](#)
- BERK, R. A. and DE LEUW, J. (1999). An evaluation of California’s inmate classification system using a generalized regression discontinuity design. *J. Amer. Statist. Assoc.* **94** 1045–1052.
- BERRY, S. M. and BERRY, D. A. (2004). Accounting for multiplicities in assessing drug safety: A three-level hierarchical mixture model. *Biometrics* **60** 418–426. [MR2066276](#)
- CATTANEO, M. D., FRANSDEN, B. and TITIUNIK, R. (2015). Randomization inference in the regression discontinuity design: An application to party advantages in the U.S. Senate. *Journal of Causal Inference* **3** 1–24.
- CELLINI, S. R., FERREIRA, F. and ROTHSTEIN, J. (2010). The value of school facility investments: Evidence from a dynamic regression discontinuity design. *Q. J. Econ.* **125** 215–261.
- CHIB, S. and GREENBERG, E. (2014). Nonparametric Bayes analysis of the sharp and fuzzy regression discontinuity designs. Technical report, Washington Univ. St Louis, Olin School of Business.
- CHIB, S. and JACOBI, L. (2015). Bayesian fuzzy regression discontinuity analysis and returns to compulsory schooling. *J. Appl. Econometrics*. Published online in Wiley Online Library (wileyonlinelibrary.com), DOI:10.1002/jae.2481.
- COOK, T. D. (2008). “Waiting for life to arrive”: A history of the regression-discontinuity design in psychology, statistics and economics. *J. Econometrics* **142** 636–654. [MR2416822](#)
- DE FINETTI, B. (1937). La prévision: ses lois logiques, ses sources subjectives. *Ann. Inst. Henri Poincaré* **7** 1–68.

- DINARDO, J. and LEE, D. S. (2011). Program evaluation and research designs. In *Handbook of Labor Economics* **4A** 463–536. Elsevier, Philadelphia, PA.
- ELLIOTT, M. R., RAGHUNATHAN, T. E. and LI, Y. (2010). Bayesian inference for causal mediation effects using principal stratification with dichotomous mediators and outcomes. *Biostatistics* **11** 353–372.
- FRANGAKIS, C. E. and RUBIN, D. B. (2002). Principal stratification in causal inference. *Biometrics* **58** 21–29. [MR1891039](#)
- FRUMENTO, P., MEALLI, F., PACINI, B. and RUBIN, D. B. (2012). Evaluating the effect of training on wages in the presence of noncompliance, nonemployment, and missing outcome data. *J. Amer. Statist. Assoc.* **107** 450–466. [MR2980057](#)
- GARIBALDI, P., GIAVAZZI, F., ICHINO, A. and RETTORE, E. (2012). College cost and time to complete a degree: Evidence from tuition discontinuities. *Rev. Econ. Stat.* **94** 699–711.
- GELMAN, A., MENG, X.-L. and STERN, H. (1996). Posterior predictive assessment of model fitness via realized discrepancies. *Statist. Sinica* **6** 733–807. [MR1422404](#)
- GELMAN, A. E. and RUBIN, D. B. (1992). Inference from iterative simulation using multiple sequences. *Statist. Sci.* **7** 457–472.
- GOSSELIN, F. (2011). A new calibrated Bayesian internal goodness-of-fit method: Sampled posterior p -values as simple and general p -values that allow double use of the data. *PLoS ONE* **6** 1–10.
- HAHN, J., TODD, P. E. and VAN DER KLAUW, W. (2001). Identification and estimation of treatment effects with a regression-discontinuity design. *Econometrica* **69** 201–209.
- IMBENS, G. W. (2004). Nonparametric estimation of average treatment effects under exogeneity: A review. *The Review of Economics and Statistics* **86** 4–29.
- IMBENS, G. W. and ANGRIST, J. D. (1994). Identification and estimation of local average treatment effects. *Econometrica* **62** 467–476.
- IMBENS, G. and KALYANARAMAN, K. (2012). Optimal bandwidth choice for the regression discontinuity estimator. *Rev. Econ. Stud.* **79** 933–959. [MR2986387](#)
- IMBENS, G. W. and LEMIEUX, T. (2008). Regression discontinuity designs: A guide to practice. *J. Econometrics* **142** 615–635. [MR2416821](#)
- IMBENS, G. W. and RUBIN, D. B. (1997). Bayesian inference for causal effects in randomized experiments with noncompliance. *Ann. Statist.* **25** 305–327. [MR1429927](#)
- IMBENS, G. W. and RUBIN, D. B. (2015). *Causal Inference—for Statistics, Social, and Biomedical Sciences: An Introduction*. Cambridge Univ. Press, New York. [MR3309951](#)
- IMBENS, G. W. and ZAJONC, T. (2011). Regression discontinuity design with multiple forcing variables. Technical report, Harvard Univ., Dept. Economics.
- JOHNSON, V. E. (2007). Bayesian model assessment using pivotal quantities. *Bayesian Anal.* **2** 719–733. [MR2361972](#)
- LEE, D. S. (2008). Randomized experiments from non-random selection in U.S. House elections. *J. Econometrics* **142** 675–697. [MR2416824](#)
- LEE, D. S. and LEMIEUX, T. (2010). Regression discontinuity designs in economics. *J. Econ. Lit.* **485** 281–355.
- LI, F., MATTEI, A. and MEALLI, F. (2015). Supplement to “Evaluating the causal effect of university grants on student dropout: Evidence from a regression discontinuity design using principal stratification.” DOI:[10.1214/15-AOAS881SUPP](#).
- LUDWIG, J. and MILLER, D. L. (2007). Does head start improve children’s life chances? Evidence from a regression discontinuity design. *Q. J. Econ.* **122** 15981–208.
- MATTEI, A., LI, F. and MEALLI, F. (2013). Exploiting multiple outcomes in Bayesian principal stratification analysis with application to the evaluation of a job training program. *Ann. Appl. Stat.* **7** 2336–2360. [MR3161725](#)
- MEALLI, F. and PACINI, B. (2013). Using secondary outcomes to sharpen inference in randomized experiments with noncompliance. *J. Amer. Statist. Assoc.* **108** 1120–1131. [MR3174688](#)

- MEALLI, F. and RAMPICHINI, C. (2012). Evaluating the effects of university grants by using regression discontinuity designs. *J. Roy. Statist. Soc. Ser. A* **175** 775–798. [MR2948374](#)
- MEALLI, F. and RUBIN, D. B. (2002). Assumptions when analyzing randomized experiments with noncompliance and missing outcomes. *Health Serv. Outcomes Res. Methodol.* **3** 225–232.
- MERCATANTI, A. (2013). A likelihood-based analysis for relaxing the exclusion restriction in randomized experiments with noncompliance. *Aust. N. Z. J. Stat.* **55** 129–153. [MR3079024](#)
- MERCATANTI, A., LI, F. and MEALLI, F. (2015). Improving inference of Gaussian mixtures using auxiliary variables. *Stat. Anal. Data Min.* **8** 34–48. [MR3315979](#)
- MURPHY, S. A. (2003). Optimal dynamic treatment regimes. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **65** 331–366. [MR1983752](#)
- RUBIN, D. B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology* **66** 688–701.
- RUBIN, D. B. (1978). Bayesian inference for causal effects: The role of randomization. *Ann. Statist.* **6** 34–58. [MR0472152](#)
- RUBIN, D. B. (1980). Discussion of “Randomization analysis of experimental data: The Fisher randomization test” by D. Basu. *J. Amer. Statist. Assoc.* **75** 591–593.
- SALES, A. and HANSEN, B. (2014). Limitless regression discontinuity: Causal inference for a population surrounding a threshold. Available at [arXiv:1403.5478](#).
- SCHWARTZ, S. L., LI, F. and MEALLI, F. (2011). A Bayesian semiparametric approach to intermediate variables in causal inference. *J. Amer. Statist. Assoc.* **31** 949–962.
- SCOTT, J. G. and BERGER, J. O. (2006). An exploration of aspects of Bayesian multiple testing. *J. Statist. Plann. Inference* **136** 2144–2162. [MR2235051](#)
- THISTLETHWAITE, D. and CAMPBELL, D. (1960). Regression-discontinuity analysis: An alternative to the ex-post facto experiment. *J. Educ. Psychol.* **51** 309–317.
- VAN DER KLAUW, W. (2002). Estimating the effect of financial aid offers on college enrollment: A regression-discontinuity approach. *Internat. Econom. Rev.* **43** 1249–1287.
- VAN DER KLAUW, W. (2008). Regression-discontinuity analysis: A survey of recent development in economics. *Labour* **22** 219–245.
- ZAJONC, T. (2012). Bayesian inference for dynamic treatment regimes: Mobility, equity, and efficiency in student tracking. *J. Amer. Statist. Assoc.* **107** 80–92. [MR2949343](#)

F. LI
DEPARTMENT OF STATISTICAL SCIENCE
DUKE UNIVERSITY
122 OLD CHEMISTRY BUILDING
DURHAM, NORTH CAROLINA 27708-0251
USA
E-MAIL: fli@stat.duke.edu

A. MATTEI
F. MEALLI
DEPARTMENT OF STATISTICS,
INFORMATICS, APPLICATIONS
UNIVERSITY OF FLORENCE
VIALE MORGAGNI, 59
50134 FLORENCE
ITALY
E-MAIL: mattei@disia.unifi.it
mealli@disia.unifi.it