



## LARF: Instrumental Variable Estimation of Causal Effects through Local Average Response Functions

Weihua An  
Indiana University

Xuefu Wang  
Indiana University

---

### Abstract

**LARF** is an R package that provides instrumental variable estimation of treatment effects when both the endogenous treatment and its instrument (i.e., the treatment inducement) are binary. The method (Abadie 2003) involves two steps. First, pseudo-weights are constructed from the probability of receiving the treatment inducement. By default **LARF** estimates the probability by a probit regression. It also provides semiparametric power series estimation of the probability and allows users to employ other external methods to estimate the probability. Second, the pseudo-weights are used to estimate the local average response function conditional on treatment and covariates. **LARF** provides both least squares and maximum likelihood estimates of the conditional treatment effects.

*Keywords:* instrumental variable, causal inference, compliers, local average response function.

---

## 1. Introduction

We created the package **LARF** (An and Wang 2015) in R (R Core Team 2016) to implement the method proposed by Abadie (2003) that provides instrumental variable (IV) estimation of causal effects when both the endogenous treatment and its instrument are binary. The main strength of the method is that it can address treatment selection without the imposition of parametric identifying assumptions.

The method is based on the potential outcomes framework (Rubin 1974, 1977). Suppose we are interested in estimating the average effect of a binary treatment  $D$  on outcome  $Y$ , namely,  $E[Y_1 - Y_0]$ . The term  $Y_1$  represents the potential outcome if an individual receives the treatment while  $Y_0$  the potential outcome if the individual does not receive the treatment. The average treatment effect (ATE) can also be viewed as a weighted average:  $\pi E[Y_1 - Y_0 | D = 1] + (1 - \pi) E[Y_1 - Y_0 | D = 0]$ , where  $\pi$  and  $1 - \pi$  are the proportions of the treated and the control units in the population, respectively.  $E[Y_1 - Y_0 | D = 1]$  represents the average

treatment effect on the treated (ATT) and  $E[Y_1 - Y_0|D = 0]$  the average treatment effect on the control (ATC). In a naive way, we may use the difference between the average outcomes of the treated group and the control group (i.e.,  $E[Y_1|D = 1] - E[Y_0|D = 0]$ ) to estimate the ATE. However, it is shown that this naive estimator is biased (Morgan and Winship 2014).

$$\begin{aligned} E[Y_1|D = 1] - E[Y_0|D = 0] - E[Y_1 - Y_0] \\ = E[Y_0|D = 1] - E[Y_0|D = 0] + (1 - \pi)(ATT - ATC). \end{aligned}$$

The bias can arise from two sources: outcome differences across the treatment groups in the absence of treatment and differential treatment effects for the treated and the control. The bias tends to be present if units which expect to benefit from the treatment select to take the treatment while units which expect not to benefit or benefit less reject the treatment. Similarly, we can show that this estimator is biased for estimating either ATT or ATC.

When there is a randomized treatment inducer  $Z$ , which is imperfectly correlated with the treatment intake  $D$ , Imbens and Angrist (1994) and Angrist, Imbens, and Rubin (1996) show that under mild assumptions, local average treatment effect (LATE) for compliers (i.e., individuals for whom  $D = 1$  if  $Z = 1$ , and  $D = 0$  if  $Z = 0$ ) can be identified by instrumenting the treatment intake  $D$  by the treatment inducer  $Z$ . Abadie (2003) extended the previous methods by allowing the treatment inducer to be randomized conditionally on the covariates and by allowing the outcome to depend on the covariates besides the treatment intake. Abadie (2003) also provided semiparametric estimations of the probability of receiving the treatment inducement, which helps to identify the treatment effects in a more robust way.

Specifically, when the treatment inducer  $Z$  is as good as randomized after conditioning on covariates  $X$ , Abadie (2003) proposed a two-stage procedure to estimate treatment effects. In the first step, it estimates the probability of receiving the treatment inducement  $P(Z = 1|X)$  in order to provide a set of pseudo-weights. Second, the pseudo-weights are used to estimate the local average response function (LARF) of the outcome conditional on the treatment and covariates. The estimated coefficient for the treatment intake  $D$  reflects the conditional treatment effect.

The main function `larf` in the **LARF** package implements the method proposed in Abadie (2003). The package is available from the Comprehensive R Archive Network (CRAN) at <https://CRAN.R-project.org/package=LARF>.

This paper proceeds as follows. In Section 2, we sketch the statistical theories behind LARF. In Section 3, we demonstrate the usage of **LARF** with both simulated and empirical examples. In Section 4, we compare our package with similar programs.

## 2. Estimation of treatment effects through LARF

### 2.1. Identification of treatment effect for compliers

Suppose there is a binary treatment inducer  $Z$ , which is as good as randomized after conditioning on the covariates. Using  $D_z$  to indicate the potential treatment intake when  $Z = z$ , we can divide the units into four groups. The first group contains compliers, i.e., those who will comply with their treatment inducement. If they are induced to take the treatment (i.e.,  $z = 1$ ), they will take the treatment (i.e.,  $D_1 = 1$ ). If they are induced into the control

Treatment inducement $Z$		Treatment intake $D$	
		0	1
0	(1) Compliers, Never Takers	(2) Always Takers, Defiers	
1	(3) Never Takers, Defiers	(4) Compliers, Always Takers	

Table 1: Dividing units into four groups by treatment inducement and intake.

condition (i.e.,  $z = 0$ ), they will not take the treatment (i.e.,  $D_0 = 0$ ). In other words, for this group,  $D_1 > D_0$ . The second group is always-takers who will take the treatment regardless of the treatment inducement, i.e.,  $D_1 = D_0 = 1$ . The third group is never-takers who will not take the treatment whatsoever, i.e.,  $D_1 = D_0 = 0$ . The last group is defiers who will do the opposite to their treatment inducement. If they are induced to take the treatment, they will opt out. But if they are induced not to take the treatment, they will opt in. Thus for this group,  $D_1 < D_0$ . Table 1 shows the breakdown of the four groups.

Based on Angrist *et al.* (1996), Abadie (2003) listed a set of assumptions for the proposed IV method to identify treatment effects. With slight adaption we present them below.

- i) Conditional ignorability:  $(Y_{dz}, D_z) \perp Z | X$ , namely, conditioning on  $X$ , the potential outcomes  $Y$  and the potential treatment intake  $D$  are independent of  $Z$ . In other words,  $Z$  is as good as randomly assigned conditional on covariates  $X$ .
- ii) Exclusion of the instrument:  $P(Y_{d1} | D = d, X = x) = P(Y_{d0} | D = d, X = x)$ . This means that  $Z$  has no direct effects on  $Y$  once actual treatment intake and covariates are controlled for. The first two assumptions ensure the exogeneity of  $Z$ .
- iii) Relevance of the instrument:  $\text{COV}(D, Z) \neq 0$ , namely,  $Z$  is predictive of  $D$ . The first three assumptions ensure that  $Z$  can serve as a valid instrument for  $D$ .
- iv) Common support:  $0 < P(Z = 1 | X) < 1$ , namely, not all units with certain covariates are induced to take or not to take the treatment.
- v) Monotonicity:  $P(D_1 \geq D_0 | X) = 1$ . This assumption ensures there are no defiers.<sup>1</sup>

Given the monotonicity assumption, we can identify the proportion of compliers, always-takers and never-takers, respectively.

$$\pi_{\text{compliers}} : P(D_1 > D_0 | X) = E[D | X, Z = 1] - E[D | X, Z = 0], \quad (1)$$

$$\pi_{\text{always-takers}} : P(D_1 = D_0 = 1 | X) = E[D | X, Z = 0], \quad (2)$$

$$\pi_{\text{never-takers}} : P(D_1 = D_0 = 0 | X) = 1 - E[D | X, Z = 1]. \quad (3)$$

Let  $g$  be a measurable real function of  $(Y, D, X)$  such that  $E|g(Y, D, X)| < \infty$ . We can show that the expectation of  $g$  is a weighted sum of the expectations of  $g$  in the three groups.

$$E[g | X] = E[g | X, D_1 > D_0]P(D_1 > D_0 | X) + E[g | X, D_1 = D_0 = 1]P(D_1 = D_0 = 1 | X) + E[g | X, D_1 = D_0 = 0]P(D_1 = D_0 = 0 | X). \quad (4)$$

<sup>1</sup>Although the monotonicity assumption may hold generally, it needs to be assessed case by case. It is worth noting that some prior work has attempted to relax this assumption, for examples, Shepherd, Redman, and Ankerst (2008) and Small, Tan, Lorch, and Brookhart (2014).

Rearranging the terms in (4) and plugging in the group proportions, [Abadie \(2003\)](#) shows

$$\mathbb{E}[g|X, D_1 > D_0] = \frac{1}{\mathbb{P}(D_1 > D_0|X)} \mathbb{E} \left[ g \left( 1 - \frac{D(1-Z)}{\mathbb{P}(Z=0|X)} - \frac{Z(1-D)}{\mathbb{P}(Z=1|X)} \right) \middle| X \right]. \quad (5)$$

Integrating out  $X$  (see page 258 of [Abadie 2003](#)), we get

$$\begin{aligned} \mathbb{E}[g|D_1 > D_0] &= \frac{1}{\mathbb{P}(D_1 > D_0)} \mathbb{E} \left[ g \left( 1 - \frac{D(1-Z)}{\mathbb{P}(Z=0|X)} - \frac{Z(1-D)}{\mathbb{P}(Z=1|X)} \right) \right] \\ &= \frac{1}{\mathbb{P}(D_1 > D_0)} \mathbb{E}[g\kappa], \end{aligned}$$

where

$$\kappa = 1 - \frac{D(1-Z)}{\mathbb{P}(Z=0|X)} - \frac{(1-D)Z}{\mathbb{P}(Z=1|X)}.$$

The results shows that, the joint distribution of  $(Y, D, X)$  is identifiable for compliers if we use the  $\kappa$  as a set of pseudo weights. Since for compliers there is no treatment selection problem, we can also identify the treatment effect of  $D$  on  $Y$ .

## 2.2. Estimation of the pseudo weights

In order to estimate the pseudo weights  $\kappa$ , it is necessary to estimate  $\mathbb{P}(Z=1|X)$ , i.e., the probability of receiving the treatment inducement. [Abadie \(2003\)](#) proposed two methods to estimate  $\mathbb{P}(Z=1|X)$ . One is based on a probit model, namely,  $\mathbb{P}(Z=1|X) = \Phi(X^\top \gamma)$ . The predicted probability along with  $D$  are then used to calculate  $\kappa$ . The probability  $\mathbb{P}(Z=1|X)$  can also be estimated by a semiparametric power series method ([Abadie 2003](#), p. 243; [Mercatanti and Li 2014](#)). In brief, the semiparametric power series method uses the power series of the covariates to predict the probability of receiving the treatment inducement. The optimal order of the power series can be determined by cross-validations. Relatively speaking, the power series method is more robust to potential model specification errors. In the package, we implement both the probit and the power series methods. See Section 3 for examples.

## 2.3. Estimation of the local average response function

Define  $\mathbb{E}[Y|X, D, D_1 > D_0]$  as the local average response function for the compliers (LARF). By construction, the results for  $g$  shown above are also applicable to the LARF. [Abadie \(2003\)](#) proposed two methods to estimate the LARF: One is based on the weighted least squares method (LS) and the other based on maximum likelihood (ML) estimation.

Suppose  $\mathbb{E}[Y|X, D, D_1 > D_0] = h(D, X; \theta)$ . Given the results in Section 2.1 and the availability of  $\kappa$ , the parameter vector  $\theta_0$  in the LS estimator is

$$\theta_0 = \underset{\theta \in \Theta}{\operatorname{argmin}} \mathbb{E}[\kappa (Y - h(D, X; \theta))^2]. \quad (6)$$

Let  $A = (D, X)$  and  $\theta = (\alpha; \beta)$ . In **LARF**, for continuous outcomes  $h = A^\top \theta = (\alpha D + X^\top \beta)$ . For binary outcomes,  $h = \Phi(A^\top \theta) = \Phi(\alpha D + X^\top \beta)$ .

**LARF** also provides ML estimates. For continuous outcomes,  $Y - (\alpha D + X^\top \beta)$  is assumed to follow a Normal distribution with a variance  $\sigma^2$ . Then

$$\theta_0 = \operatorname{argmax}_{\theta \in \Theta} \mathbb{E}[\kappa \{-\ln(\sigma) - \frac{1}{2\sigma^2}(Y - (\alpha D + X^\top \beta))^2\}]. \quad (7)$$

For binary outcomes,  $Y$  is approximated by a probit function. Then we have

$$\theta_0 = \operatorname{argmax}_{\theta \in \Theta} \mathbb{E}[\kappa \{Y \ln \Phi(\alpha D + X^\top \beta) + (1 - Y) \ln \Phi(-\alpha D - X^\top \beta)\}]. \quad (8)$$

The variances of the coefficients in **LARF** are estimated according to Theorem 4.2 in [Abadie \(2003\)](#) if the probability of receiving treatment inducement is estimated by a probit model. The variances are estimated according to Theorem 4.5 in [Abadie \(2003\)](#) if the probability of receiving treatment inducement is estimated by the power series method.

### 3. Package overview and examples

#### 3.1. Usage of the **LARF**

The `larf` function in **LARF** implements the aforementioned methods. Its basic usage is:

```
larf(formula, treatment, instrument, data, method = "LS", AME = FALSE,
     optimizer = "Nelder-Mead", zProb = NULL)
```

where the arguments are defined as follows.

- **formula**: specification of the outcome model in the form like either  $y \sim x_1 + x_2$  or  $y \sim X$  where  $X$  is a matrix containing all the covariates excluding the treatment. Also support multi-part formulas ([Zeileis and Croissant 2010](#)). For example,  $y + d \sim x_1 + x_2 \mid z$ , where  $d$  represents the treatment and  $z$  the instrument.
- **treatment**: a vector containing the binary treatment.
- **instrument**: a vector containing the binary instrument for the endogenous treatment.
- **data**: an optional data frame. If unspecified, the data will be taken from the working environment.
- **method**: the estimation method to be used. The default is "LS", standing for least squares. "ML", standing for maximum likelihood, is an alternative.
- **AME**: whether average marginal effects (AME) should be reported. The default is **FALSE**, in which case marginal effects at the means (MEM) are reported.
- **optimizer**: the optimization algorithm for the ML method. It should be one of "Nelder-Mead", "BFGS", "CG", "L-BFGS-B", "SANN", or "Brent". See `optim` in R for more detail.
- **zProb**: a vector containing the probability of receiving the treatment inducement that has been estimated externally.

The `larf` function returns the estimated coefficients for the treatment and covariates and their standard errors. For binary outcomes, it also returns the marginal effects at the means (MEM) or the average marginal effects (AME) if requested.<sup>2</sup> `larf` is the high-level interface to the work-horse function `larf.fit`. A set of standard methods (including `print`, `summary`, `vcov`, and `predict`) can be used to extract the corresponding information from a ‘`larf`’ object. For more detail, please see `?larf.fit`.

### 3.2. Simulations

We conduct simulations to show the properties of the LARF method. The equations we used to generate the data are as follows.

$$Y^* = \alpha D + X^\top \beta + e, \quad (9)$$

$$D^* = (Z, X)^\top \gamma + v. \quad (10)$$

The covariates include  $X_1$  and  $X_2$  and are generated according to

$$X_{1i} \sim \text{Normal}(0, 1),$$

$$X_{2i} \sim \text{Normal}(0, 2).$$

The treatment inducer  $Z$  is generated according to

$$Z_i \sim \text{Bernoulli}(0.5).$$

The error terms are correlated across equations and are generated according to

$$\begin{pmatrix} e_i \\ v_i \end{pmatrix} \sim \text{Normal}\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & 0.8 \\ 0.8 & 1 \end{pmatrix}\right).$$

We then generate  $D_i^* = -0.1 + 0.7Z_i + 0.3X_{1i} - 0.5X_{2i} + v_i$ . We view  $D_i^*$  as the latent probability of receiving the treatment. Corresponding to a probit model, the observed treatment intake  $D_i$  is then set to one if  $D_i^*$  is larger than zero and to zero otherwise.<sup>3</sup> We generate the latent outcome  $Y_i^* = -0.2 + 0.8D + 0.4X_{1i} - 0.6X_{2i} + e_i$ . Similarly, we set the observed outcome  $Y_i$  to one if  $Y_i^*$  is larger than zero and to zero otherwise.

The correlation in the error terms leads the treatment  $D$  to be correlated with  $e$ . Thus a simple regression of the outcome on  $D$  and the covariates will lead to biased estimates of the treatment effects. Since  $Z$  is randomly assigned and is uncorrelated with the outcome except via its correlation with the treatment, it can be used as an instrument for  $D$  to estimate the treatment effect.

---

<sup>2</sup>Marginal effect at the means for a continuous variable is the partial derivative of the outcome with respect to the variable evaluated at the means of other variables, i.e.,  $\phi(\bar{A}^\top \hat{\theta})\hat{\theta}$ . For a binary variable  $C$ , it is the change in the probability when the variable is changed from zero to one while holding other variables at their means, i.e.,  $\Phi(\bar{A}^\top \hat{\theta}|C = 1) - \Phi(\bar{A}^\top \hat{\theta}|C = 0)$ . Average marginal effect is the average of individual marginal effects, i.e.,  $E[\phi(A^\top \hat{\theta})]\hat{\theta}$ . Average marginal effect for a binary variable is the average change in the probability when the variable is changed from zero to one, i.e.,  $E[\Phi(A^\top \hat{\theta}|C = 1) - \Phi(A^\top \hat{\theta}|C = 0)]$ . These marginal effects are widely used in social sciences. But it is worth noting that the term ‘‘marginal effects’’ has been used differently, for example, to refer to the treatment effects estimated by weighting regressions with the inverse probability of treatment (Bang and Robins 2005; Lunceford and Davidian 2004).

<sup>3</sup>This is the latent variable approach to motivating the probit model, where  $P(D_i^* > 0) = P(D_i = 1)$ .

	Naive		LARF		2SLS	
Outcome	Bias	MSE	Bias	MSE	Bias	MSE
Continuous	1.25	1.56	-0.02	0.04	-0.02	0.04
Binary	1.78	3.19	0.52	0.30	-0.37	0.14

Table 2: Results of the simulations. Note: Rows 1 and 2 show the simulation results for continuous outcomes and binary outcomes, respectively. Bias is the 10% trimmed mean of (estimates  $-$  0.8). Namely, 10% of very large or very small estimates are removed before calculating the mean deviation in order to reduce the influence of outliers. MSE is the mean-square-error. Columns (1) and (2) show the results when the outcomes are directly predicted from the treatment and covariates. Columns (3) and (4) show the results of the LARF estimates. Columns (5) and (6) show the results of the 2SLS estimates.

We conduct two sets of simulations. First, we treat the latent outcome  $Y^*$  as a observed continuous outcome. Second, we treat  $Y$  as the observed binary outcome. For comparison, we also provide naive estimates based on ordinary least squares (OLS) for the continuous outcome and probit for the binary outcome and corresponding estimates based on two-stage least squares (2SLS). We set the sample size at 2,000 and conduct 500 rounds of simulations for each type of outcome. The results are shown in Table 2.

When the outcome is continuous, the naive estimation has both much larger bias and mean-squared-error (MSE) than the LARF and the 2SLS estimations. Between LARF and 2SLS estimations, there are no discernible differences. When the outcome is binary, the naive estimation also has much larger bias and MSE than LARF and 2SLS. There is also little evidence that one is qualitatively better than the other between LARF and 2SLS. This happens maybe because we have a constant treatment effect. If there is effect heterogeneity, LARF may perform better than 2SLS. In addition, LARF allows covariates to have nonlinear effects on the outcome, which is an attractive feature in many contexts. Although the above results are based on a specific simulation, the patterns of the results should hold generally because the simulation has a general structure.

### 3.3. An empirical example

We show an empirical example using the data on eligibility for and participation in 401(k) from the Wooldridge data sets (Wooldridge 2010).<sup>4</sup> The basic research questions are to what extent participation in 401(k) – usually an employer-sponsored retirement saving plan, affects individuals’ net family financial assets and participation in other retirement saving plans (e.g., Individual Retirement Accounts (IRA)). Because individuals may selectively participate in 401(k) due to their different unobserved preferences for savings, the naive estimate of the treatment effect by comparing the average outcomes between the participants and non-participants of 401(k) may be biased.

The data includes both a continuous outcome on net family financial assets (`nettfa`) and a binary outcome on participation in IRA (`pira`). The treatment is participation in 401(k) (`p401k`). Eligibility for 401(k), `e401k`, is used as an instrument for participation in 401(k).

<sup>4</sup>The data is entitled `401ksubs.dta` in Stata (StataCorp. 2015) format, available at <http://www.stata.com/texts/eacsap/>. Based on the summary statistics of the data, it appears to be the same data used in Poterba, Venti, and Wise (1995), Poterba, Venti, and Wise (1997), and Abadie (2003).

Variables	Entire sample		Participants		Non-participants		<i>p</i>
	Mean	SD	Mean	SD	Mean	SD	
Participation in 401(k) (binary)	0.28	0.45					
Eligibility for 401(k) (binary)	0.39	0.49	1.00	0.00	0.16	0.37	0.00
Participation in IRA (binary)	0.25	0.44	0.36	0.48	0.21	0.41	0.00
Married (binary)	0.63	0.48	0.70	0.46	0.60	0.49	0.00
Family net financial assets (\$1000)	19.07	63.96	38.47	79.27	11.67	55.29	0.00
Family income (\$1000)	39.25	24.09	49.82	26.81	35.22	21.65	0.00
Age (continuous)	41.08	10.30	41.51	9.65	40.91	10.53	0.01
Family Size (count)	2.89	1.53	2.92	1.47	2.87	1.55	0.17
<i>N</i>	9275		6713		2562		

Table 3: Summary statistics of the data. Note: The data is obtained from the Wooldridge data sets (Wooldridge 2010), available at <http://www.stata.com/texts/eacsap/>. The type of data is shown after each variable’s name. The first four variables are binary. Thus their means can be interpreted as proportions. For example, the mean of “Married” is 0.63, which indicates 63% of the subjects are married. Its standard deviation equals 0.48, i.e., the square root of  $0.37 \times (1 - 0.37)$ . Participants are those who have participated in 401(k) while non-participants are those who have not enrolled in 401(k). The last column shows the *p* values of the *t*-tests that test whether the means of the covariates between the participants and non-participants are equal.

Since an individual has to be eligible for 401(k) to participate in it, there are neither defiers nor always-takers in this case. Because 401(k) eligibility is determined by employers, it is likely independent of employees’ unobserved preferences for savings. To what extent 401(k) eligibility can legitimately serve as an instrument for participation in 401(k), however, is debatable (Engen, Gale, and Scholz 1994; Poterba *et al.* 1995). But for illustrative purposes, following Abadie (2003), we assume that it can be used as an instrument for participation in 401(k). Table 3 presents the summary statistics of the data.<sup>5</sup>

To demonstrate the usage of `larf`, we load the package **LARF** and the data `c401k`.

```
R> library("LARF")
R> data("c401k", package = "LARF")
R> attach(c401k)
```

We first estimate the effect of participation in 401(k) on the net family financial assets (`nettfa`, continuous). The covariates include an intercept, family income, age, age<sup>2</sup>, marital status, and family size. We instrument participation in 401(k) by 401(k) eligibility.

```
R> est1 <- larf(nettfa ~ inc + age + agesq + marr + fsize, treatment = p401k,
+ instrument = e401k, data = c401k)
R> summary(est1)
```

Call:

```
larf(formula = nettfa ~ inc + age + agesq + marr + fsize, treatment = p401k,
```

<sup>5</sup>Similar summary statistics for the same data can be found in Table 1 of Abadie (2003).

```
instrument = e401k, data = c401k)
```

	Estimate	SE	P
Treatment	9.4947	2.1582	0.0000
(Intercept)	-7.6793	14.8525	0.6051
inc	1.0130	0.0908	0.0000
age	-1.4790	0.8047	0.0661
agesq	0.0308	0.0099	0.0020
marr	-7.4821	2.5607	0.0035
fsize	-1.0573	0.6262	0.0913

The results show that holding everything else constant, participation in 401(k) increases family net financial assets by about 9.49 thousand dollars on average (SE = 2.16 \$1,000,  $p < 0.001$ ). By default the above calculation estimates the probability of being eligible for 401(k) using a probit model. To check the robustness of the results, we also estimate the probability using a semiparametric power series method (Abadie 2003; Mercatanti and Li 2014). We set the maximal order of the covariates power series to be 5. After running `npse` we find the optimal order of power series is 2 based on cross-validations.

```
R> ps <- npse(e401k ~ inc + age + agesq + marr + fsize, order = 5)
R> ps$Lambda
```

```
[1] 2
```

Then we use the optimal order of covariates power series to predict the probability of being eligible for 401(k) and use the predict probability to estimate the treatment effect in `larf`.

```
R> est2 <- larf(netffa ~ inc + age + agesq + marr + fsize, treatment = p401k,
+ instrument = e401k, data = c401k, zProb = ps$fitted)
R> summary(est2)
```

Call:

```
larf(formula = netffa ~ inc + age + agesq + marr + fsize, treatment = p401k,
      instrument = e401k, data = c401k, zProb = ps$fitted)
```

	Estimate	SE	P
Treatment	12.3388	5.9162	0.0370
(Intercept)	-26.3848	52.7782	0.6171
inc	0.8905	0.3174	0.0050
age	-0.4038	2.9795	0.8922
agesq	0.0180	0.0358	0.6151
marr	-5.5235	6.0013	0.3574
fsize	-1.3495	1.7367	0.4371

The results show that holding everything else constant, participation in 401(k) increases net family financial assets by about 12.34 thousand dollars on average (SE = 5.92 \$1,000,  $p < 0.05$ ). Hence, the main finding seems to be quite robust.

We then estimate the effect of participating in 401(k) on the possession of an IRA account.

```
R> est3 <- larf(pira ~ inc + age + agesq + marr + fsize, p401k, e401k,
+ data = c401k)
R> summary(est3)
```

Call:

```
larf(formula = pira ~ inc + age + agesq + marr + fsize, treatment = p401k,
      instrument = e401k, data = c401k, AME = FALSE)
```

	Estimate	SE	P	MEM	MEM-SE	MEM-P
Treatment	0.0788	0.0501	0.1157	0.0254	0.0161	0.1144
(Intercept)	-3.5598	0.4673	0.0000	-0.7428	0.0092	0.0000
inc	0.0202	0.0012	0.0000	0.0065	0.0004	0.0000
age	0.0801	0.0225	0.0004	0.0257	0.0072	0.0003
agesq	-0.0005	0.0003	0.0351	-0.0002	0.0001	0.0342
marr	0.1292	0.0701	0.0654	0.0410	0.0219	0.0614
fsize	-0.1268	0.0231	0.0000	-0.0407	0.0074	0.0000

The first two columns show the estimated coefficients and standard errors. Columns (4) and (5) show the marginal effects at the means and their standard errors. The results suggest that participation in 401(k) increases the probability of participation in IRA by 0.025 (SE = 0.016). But the effect is not statistically significant at the 5% level.<sup>6</sup> We also estimate the average marginal effects. Again, there is no significant evidence that participation in 401(k) crowds out IRA participation.

```
R> est4 <- larf(pira ~ inc + age + agesq + marr + fsize, p401k, e401k,
+ data = c401k, AME = TRUE)
R> summary(est4)
```

Call:

```
larf(formula = pira ~ inc + age + agesq + marr + fsize, treatment = p401k,
      instrument = e401k, data = c401k, AME = TRUE)
```

	Estimate	SE	P	AME	AME-SE	AME-P
Treatment	0.0788	0.0501	0.1157	0.0218	0.0139	0.1167
(Intercept)	-3.5598	0.4673	0.0000	-0.7171	0.0104	0.0000
inc	0.0202	0.0012	0.0000	0.0055	0.0003	0.0000
age	0.0801	0.0225	0.0004	0.0219	0.0061	0.0004
agesq	-0.0005	0.0003	0.0351	-0.0001	0.0001	0.0348
marr	0.1292	0.0701	0.0654	0.0352	0.0190	0.0641
fsize	-0.1268	0.0231	0.0000	-0.0347	0.0062	0.0000

We also estimate the effects of participation in 401(k) on family net financial assets by OLS and 2SLS and on IRA participation by probit, 2SLS, IV probit, and bivariate probit. The

---

<sup>6</sup>We confirmed that both our LARF estimates were the same as those estimated by Abadie's code in MATLAB (available at <http://www.hks.harvard.edu/fs/aabadie/larf.html>). These estimates are slightly different from those shown in Abadie (2003), probably because the latter calculated the marginal effects at the covariate means of the treated units.

A. Family net financial assets (continuous)	Coef.	SE	
OLS	13.527	1.810	***
2SLS	9.419	2.152	***
LARF	9.495	2.158	***
B. IRA participation (binary)	MEM	SE	
Probit	0.062	0.011	***
2SLS	0.027	0.013	**
IV probit	0.110	0.047	**
Bivariate probit	0.031	0.012	***
LARF	0.025	0.016	

Table 4: Effects of participation in 401(k). Note: The data is from the Wooldridge data sets (Wooldridge 2010), available at <http://www.stata.com/texts/eacsap/>. The dependent variables are net family financial assets and participation in IRA, respectively. The covariates include a constant, participation in 401(k), family income, age, age square, marital status, and family size. Participation in 401(k) is instrumented by eligibility for 401(k). Panel A shows the estimated coefficients for the treatment, which reflect the average increase in family net financial assets when a person participates in 401(k) while holding everything else constant. Panel B shows the marginal effects for the treatment while holding other covariates at their means, which reflect the change in probability of participating IRA when an average person participates in 401(k). Significance level: \*,  $p < 0.10$ ; \*\*,  $p < 0.05$ ; \*\*\*,  $p < 0.01$ .

results are shown in Table 4. For net family financial assets, both the 2SLS and the LARF estimates are much smaller than the OLS estimate, suggesting that the latter may have been biased upwardly by treatment selection. For IRA participation, the LARF estimate is the only one that is not statistically significant at the 5% level, i.e., the only one that is not in conflict with the theoretical prediction. These estimates are only for illustrative purposes, as models with different covariates will render maybe somewhat different results.

## 4. Comparison with similar programs

**LARF** fills an important gap in IV estimations. Programs like **systemfit** (Henningsen and Hamann 2015, 2007) in R and **ivregress** (StataCorp. 2015) and **ivregress2** (Wada 2012) in Stata have been widely used for IV estimations. They fit linear regressions and work well if both the outcome and the endogenous treatment are continuous. But if the dependent variable is binary, the predictions can be inappropriately outside of (0, 1).<sup>7</sup> **ivprobit** (StataCorp. 2015) in Stata works for binary outcomes, but still requires a continuous treatment. The bivariate probit model (e.g., **biprobit** (StataCorp. 2015) in Stata and **binom2.rho** of **VGAM** (Yee 2015a,b) in R) can estimate causal effects with binary endogenous treatment, but it requires the outcome to be binary and also (sometimes unrealistically) assumes that the joint distribution for the outcome and the treatment is correctly known. In short, when both the endogenous treatment and its instrument are binary, **LARF** can provide more appropriate estimations of treatment effects for both continuous and binary outcomes.

<sup>7</sup>Note that when the treatment inducer is binary and fully randomized, the predicted outcome by **ivregress** or **ivregress2** will be within the range (0, 1) (Angrist *et al.* 1996).

## References

- Abadie A (2003). “Semiparametric Instrumental Variable Estimation of Treatment Response Models.” *Journal of Econometrics*, **113**, 231–263. doi:10.1016/s0304-4076(02)00201-4.
- An W, Wang X (2015). *LARF: Local Average Response Functions for Instrumental Variable Estimation of Treatment Effects*. R package version 1.3, URL <https://CRAN.R-project.org/package=LARF>.
- Angrist J, Imbens G, Rubin D (1996). “Identification of Causal Effects Using Instrumental Variables.” *Journal of the American Statistical Association*, **91**, 444–455. doi:10.1080/01621459.1996.10476902.
- Bang H, Robins J (2005). “Doubly Robust Estimation in Missing Data and Causal Inference Models.” *Biometrics*, **61**, 962–972. doi:10.1111/j.1541-0420.2008.01025.x.
- Engen EM, Gale WG, Scholz JK (1994). “Do Saving Incentives Work?” *Brookings Papers on Economic Activity*, **1**, 85–180. doi:10.2307/2534631.
- Henningsen A, Hamann J (2015). *systemfit: Estimating Systems of Simultaneous Equations*. R package version 1.1-16, URL <https://CRAN.R-project.org/package=systemfit>.
- Henningsen A, Hamann JD (2007). “systemfit: A Package for Estimating Systems of Simultaneous Equations in R.” *Journal of Statistical Software*, **23**(4), 1–40. doi:10.18637/jss.v023.i04.
- Imbens G, Angrist J (1994). “Identification and Estimation of Local Average Treatment Effect.” *Econometrica*, **62**, 467–476. doi:10.2307/2951620.
- Lunceford JK, Davidian M (2004). “Stratification and Weighting via the Propensity Score in Estimation of Causal Treatment Effects: A Comparative Study.” *Statistics in Medicine*, **23**, 2937–2960. doi:10.1002/sim.1903.
- Mercatanti A, Li F (2014). “Do Debit Cards Increase Household Spending? Evidence from a Semiparametric Causal Analysis of a Survey.” *The Annals of Applied Statistics*, **8**(4), 2405–2508. doi:10.1214/14-aos784.
- Morgan S, Winship C (2014). *Counterfactuals and Causal Inference*. 2nd edition. Cambridge University Press.
- Poterba J, Venti S, Wise D (1995). “Do 401(k) Contributions Crowd out Other Personal Saving?” *Journal of Public Economics*, **58**, 1–32. doi:10.1016/0047-2727(94)01462-w.
- Poterba J, Venti S, Wise D (1997). “Personal Retirement Saving Programs and Asset Accumulation: Reconciling the Evidence.” *Working Paper 5599*, National Bureau of Economic Research. URL <http://www.nber.org/papers/w5599>.
- R Core Team (2016). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.
- Rubin D (1974). “Estimating Causal Effects of Treatments in Randomized and Nonrandomized Studies.” *Journal of Educational Psychology*, **66**, 688–701. doi:10.1037/h0037350.

- Rubin D (1977). “Assignment to Treatment Group on the Basis of a Covariate.” *Journal of Educational Statistics*, **2**, 1–26. doi:10.2307/1164933.
- Shepherd B, Redman M, Ankerst D (2008). “Does Finasteride Affect the Severity of Prostate Cancer? A Causal Sensitivity Analysis.” *Journal of the American Statistical Association*, **103**, 1392–1404. doi:10.1198/016214508000000706.
- Small D, Tan Z, Lorch S, Brookhart M (2014). “Instrumental Variable Estimation When Compliance Is Not Deterministic: The Stochastic Monotonicity Assumption.” *Manuscript*. <http://xxx.tau.ac.il/abs/1407.7308>.
- StataCorp (2015). *Stata Data Analysis Statistical Software: Release 14*. StataCorp LP, College Station, TX. URL <http://www.stata.com/>.
- Wada R (2012). “ivregress2: Stata Module to Export First and Second-Stage Results Similar to ivregress.” Statistical Software Components, Boston College Department of Economics. URL <http://ideas.repec.org/c/boc/bocode/s457466.html>.
- Wooldridge J (2010). *Econometric Analysis of Cross Section and Panel Data*. 2nd edition. MIT Press.
- Yee T (2015a). **VGAM**: *Vector Generalized Linear and Additive Models*. R package version 0.9-8, URL <https://CRAN.R-project.org/package=VGAM>.
- Yee TW (2015b). *Vector Generalized Linear and Additive Models: With an Implementation in R*. Springer-Verlag, New York.
- Zeileis A, Croissant Y (2010). “Extended Model Formulas in R: Multiple Parts and Multiple Responses.” *Journal of Statistical Software*, **34**(1), 1–13. doi:10.18637/jss.v034.i01.

**Affiliation:**

Weihua An  
Departments of Sociology and Statistics  
752 Ballantine Hall  
1020 East Kirkwood Avenue  
Indiana University  
Bloomington, IN 47405, United States of America  
E-mail: [weihuaan@indiana.edu](mailto:weihuaan@indiana.edu)  
URL: <http://mypage.iu.edu/~weihuaan/>