

A Service of

ZBW

Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

van den Berg, Gerard J.

Working Paper An economic analysis of exclusion restrictions for instrumental variable estimation

IZA Discussion Papers, No. 2585

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: van den Berg, Gerard J. (2007) : An economic analysis of exclusion restrictions for instrumental variable estimation, IZA Discussion Papers, No. 2585, Institute for the Study of Labor (IZA), Bonn,

https://nbn-resolving.de/urn:nbn:de:101:1-2008042593

This Version is available at: https://hdl.handle.net/10419/34130

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet. or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



WWW.ECONSTOR.EU

IZA DP No. 2585

An Economic Analysis of Exclusion Restrictions for Instrumental Variable Estimation

Gerard J. van den Berg

January 2007

Forschungsinstitut zur Zukunft der Arbeit Institute for the Study of Labor

An Economic Analysis of Exclusion Restrictions for Instrumental Variable Estimation

Gerard J. van den Berg

Princeton University, Free University Amsterdam, IFAU-Uppsala, IFS, CREST, CEPR and IZA

> Discussion Paper No. 2585 January 2007

> > IZA

P.O. Box 7240 53072 Bonn Germany

Phone: +49-228-3894-0 Fax: +49-228-3894-180 E-mail: iza@iza.org

Any opinions expressed here are those of the author(s) and not those of the institute. Research disseminated by IZA may include views on policy, but the institute itself takes no institutional policy positions.

The Institute for the Study of Labor (IZA) in Bonn is a local and virtual international research center and a place of communication between science, politics and business. IZA is an independent nonprofit company supported by Deutsche Post World Net. The center is associated with the University of Bonn and offers a stimulating research environment through its research networks, research support, and visitors and doctoral programs. IZA engages in (i) original and internationally competitive research in all fields of labor economics, (ii) development of policy concepts, and (iii) dissemination of research results and concepts to the interested public.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

IZA Discussion Paper No. 2585 January 2007

ABSTRACT

An Economic Analysis of Exclusion Restrictions for Instrumental Variable Estimation^{*}

Instrumental variable estimation requires untestable exclusion restrictions. With policy effects on individual outcomes, there is typically a time interval between the moment the agent realizes that he may be exposed to the policy and the actual exposure or the announcement of the actual treatment status. In such cases there is an incentive for the agent to acquire information on the value of the IV. This leads to violation of the exclusion restriction. We analyze this in a dynamic economic model framework. This provides a foundation of exclusion restrictions in terms of economic behavior. The results are used to describe policy evaluation settings in which instrumental variables are likely or unlikely to make sense. For the latter cases we analyze the asymptotic bias. The exclusion restriction is more likely to be violated if the outcome of interest strongly depends on interactions between the agent's effort before the outcome is realized and the actual treatment status. The bias has the same sign as this interaction effect. Violation does not causally depend on the weakness of the candidate instrument or the size of the average treatment effect. With experiments, violation is more likely if the treatment and control groups are to be of similar size. We also address side-effects. We develop a novel economic interpretation of placebo effects and provide some empirical evidence for the relevance of the analysis.

JEL Classification: C31, C21, D81, J68, D82, D83, D84, C35, C51

Keywords: treatment, policy evaluation, information, selection effects, randomization, placebo effect

Corresponding author:

Gerard J. van den Berg Department of Economics Free University of Amsterdam De Boelelaan 1105 1081 HV Amsterdam The Netherlands E-mail: gjvdberg@xs4all.nl

^{*} Thanks to Jim Heckman, Geert Ridder, Ulrich Müller, Angus Deaton, and participants in seminars at Princeton, Columbia/NYU, Harvard, and Aarhus, and conferences in London, Berlin, Paris, and Bonn, for useful suggestions. The work in this paper is part of the ESRC and IFAU Research Programs "Advancing Programme Evaluation Methods".

1 Introduction

Instrumental variable estimation has since long been a standard econometric technique for dealing with endogeneity and selection issues in general, and for non-experimental policy evaluation in particular (see e.g. Angrist, Imbens and Rubin, 1996, Heckman, LaLonde and Smith, 1999, and Blundell and MaCurdy, 1999, for surveys). Basically, if one is interested in the effect of a "treatment variable" on an outcome variable, and the treatment is not exogenously assigned, then one may perform causal inference by exploiting the presence of variables that causally affect the treatment status but do not have a direct causal effect on the outcome. The latter restriction is called an exclusion restriction. Exclusion restrictions are identifying restrictions, so they can not be tested.¹ This means that empirical results critically depend on the validity of the exclusion restriction, and that this restriction needs to be justified on a priori grounds.

With policy effects on individual outcomes, there is typically a time interval between the moment the agent realizes that he may be exposed to the policy and the actual exposure. For example, unemployed workers are aware of the existence of policies leading to treatments at some point of time in the future. As long as the instrumental variable affecting the treatment does not have a causal effect on the individual's behavior, the exclusion restriction is not violated. Often, a sufficient condition for this is that the agent does not observe the value of the instrumental variable. However, there is an incentive for the agent to acquire information on this value. After all, the probability of exposure to treatment is a determinant of the optimal strategy, and the more the agent knows about it, the better he can fine-tune his behavior in response to this, and the higher his expected present value will be. The agent's strategy affects the outcome of interest. Thus, the acquisition of the value of the variable that is used by the econometrician as instrumental variable leads to violation of the exclusion restriction and to incorrect empirical inference.²

As an example, consider participation in a job search assistance program for

 $^{^{1}}$ We abstract from the use of testable restrictions on the joint distribution of the observables, as captured by bounds for binary settings; see e.g. Robins and Greenland (1996) for a brief overview.

²Earlier studies mentioning similar arguments, tacitly assuming that acquisition is free, include Abbring and Van den Berg (2003). Modelling that agents use available information on determinants of future (policy) events that affect the outcome of interest goes back to at least the rational expectations literature; see e.g. Hansen and Sargent (1980). Note that we are not concerned with mechanical program lock-in effects that may affect the outcomes of participants before the end of the actual treatment participation.

unemployed individuals, where the policy intensity differs across two otherwise identical geographical regions. For example, in one region, the budget for the program per potential participant may be larger, so that the individual probability of being treated is larger, holding everything else constant. An individual may be aware of the distribution of policy intensities but not know his personally relevant intensity, in which cases a regional dummy indicator may be a valid instrumental variable. If the individual finds out his relevant intensity then he will typically use this information before the treatment is realized. For example, it may be optimal to reduce the job search effort more if the treatment probability is large, because it is cheaper to provide effort after the treatment. This may lead to an under-estimation of the program effect on the employment rate say one year after inflow into unemployment.

In this paper we investigate, in the context of a dynamic economic framework, under which conditions it is optimal for the agent to acquire the value of the intended instrumental variable. This provides a foundation of exclusion restrictions in terms of economic behavior that takes costs and benefits into account. The results are used to describe the policy evaluation settings in which instrumental variables are likely or unlikely to make sense. This is especially useful since by definition no empirical evidence is available on the validity of exclusion restrictions. We also analyze the asymptotic bias of the instrumental variable estimator in case of violation of the exclusion restriction.

At first sight one may think that information acquisition does not take place if and only if the acquisition costs are high, and that therefore the conclusion is simply going to be that instrumental variables estimation is particularly useful to study policy effects for agents with scarce resources. For active labor market policy analysis this would mean that it is particularly useful for agents at the bottom of the labor market, which coincides with the target group of most of these policy measures. However, this line of reasoning ignores the role of the value of the information that is acquired. We show that this leads to a different set of conclusions. The results point at the importance of the extent to which the treatment status and the agent's effort interact in the outcome.

The literature on instrumental variable estimation has recently been concerned with the use of so-called weak instruments, i.e. instrumental variables that are only weakly related to the treatment status (see e.g. Stock, Wright and Yogo, 2002, for a survey). Is is sometimes argued that weak instruments have the advantage that they are less likely to be used by agents as direct causal inputs into the outcome of interest. We argue that in certain cases this line of reasoning is incorrect, and therefore this advantage of weak instruments may have been over-estimated.

Our results can be straightforwardly applied to situations in which the treatment variable is not a policy variable. Also, the costs of information may cover not only monetary costs but also other types of effort.³

We extend the framework and the results in a number of directions. First, notice that the outcome is a function of the agent's efforts and the treatment status. In the baseline analyses, we assume that the expected utility that an agent ultimately derives from a certain combination of effort and treatment is the same as the expected value of the corresponding outcome. For example, an unemployed individual may only be concerned about the probability of finding work, which is also the outcome in which the researcher is interested. This assumption is in line with dynamic economic theories such as job search theory. However, an agent may also be concerned about side-effects of the treatment that are not reflected in the measured outcome. We therefore extend the model to allow the agent's ultimate utility to be systematically different from the corresponding outcome. This extension also covers cases in which the agent does not know the treatment effect and can only guess it, as is the case in medical trials and pilot experimental evaluations of novel labor market programs. In all of these cases, the decision whether to acquire the value of the candidate instrument is driven by the agent's utility function, whereas the magnitude of the asymptotic bias also depends on the actual effects of effort and treatment on the outcome.

Throughout the paper we use the above example about participation in a job search assistance program as the leading example. Essentially equivalent examples concern the evaluation of training programs for unemployed workers and punitive sanctions and monitoring programs for unemployment insurance recipients (where the treatment is a punitive sanction and the policy intensity of interest is the intensity with which agents' search effort is monitored). The corresponding empirical literature contains many cases in which region is used as an instrument. A similar example concerns the use of geographical distance to college as an instrument to study the returns to education. In this case, parents of children who live far from a college may provide substitute educational support.

³Of course, exclusion restrictions for instrumental variable estimation may be violated for other reasons than those considered in this paper. Notably, the agent's value of the candidate instrument may be affected by unobserved characteristics of the agent that have a direct causal effect on the outcome variable. (This is prevented if the candidate instrument is the result of a deliberate randomization, like a deliberately randomized intention-to-treat variable in a social experiment.) Rosenzweig and Wolpin (2000) consider violations in a dynamic setup that are due to events occurring between the treatment and the outcome, whereas we are primarily concerned with behavior before realization of the treatment.

Clearly, the analysis covers randomized experiments with delayed treatment as well. There is also some relevance for the evaluation of (double-)blind experiments of say a medication to treat a disease, where the randomized intention to treat equals the treatment status, the treatment is immediate, and the treatment status is supposed to be unobserved until the outcome is observed. In the case of life-threatening diseases, an individual has an incentive to find out whether he receives medication or a placebo, for example by sending one tablet to a laboratory. If he discovers that he receives a placebo then he may choose a different lifestyle, which in turn affects the outcome. Another option is to share the tablets among participating individuals. See Epstein (1996) and Schuklenk (2003) for examples concerning experiments of AIDS medication, and Schulz (1995) for a pervasive account of the high occurrence of deciphering in general. This can be translated into our model framework. Our results concerning the determinants of the asymptotic effect bias can be used to understand the (sign of the) observed association between the treatment effect estimate and the degree of methodological quality of the study (see e.g. Schulz et al., 1995, and Schulz and Grimes, 2002).

The mechanisms described above should be distinguished from non-compliance upon assignment of the treatment status, when the actual treatment status is observed by the analyst. Under certain conditions, IV estimation can accommodate selective non-compliance (Imbens and Angrist, 1994). Nevertheless, it is interesting to study the determinants of non-compliance behavior, and we show that our model framework is amenable to this purpose. In particular, we briefly consider when agents selectively choose to become non-compliers if they are able to use information on determinants of the treatment assignment process. This establishes a relation to the literature on rational non-compliance behavior in experiments (see in particular Philipson and DeSimone, 1997, and Chan and Hamilton, 2006, for elaborate dynamic learning models).

We also relate to the literature on placebo effects. We provide an economic explanation for so-called placebo effects in experiments. This does not assume the availability of information on the treatment status, but instead rests on the principle that agents rationally use information on the randomization probability to adjust their behavior towards maximization of their expected present value.

The paper is organized as follows: Section 2 presents the model framework, Section 3 derives the results concerning information acquisition, Section 4 discusses the implications for instrumental variable estimation, notably the asymptotic bias of the estimator in case of violation of the exclusion restriction. We also discuss valid testing of the null hypothesis of no causal treatment effect. Section 5 concludes.

2 The model framework

The main insights and results can be derived in a simple model framework with three time periods.

Consider an assignment process leading to the actual treatment status of an agent or decision maker.⁴ We assume that this process depends on a variable Z. At the onset of the first period, Z is realized separately for each agent. We take the distribution of Z as exogenously given and assume that Z is dispersed, so var(Z) > 0. Z is the candidate instrumental variable. Each agent knows the distribution of Z across all agents. However, an agent does not necessarily know his own realization of Z. More precisely, in Period 1, each agent decides whether to acquire information on Z at cost $\gamma > 0$ or not.

At the onset of Period 2, the agent determines his optimal strategy or effort a. The agent may or may not know his private value of Z when determining a, but we assume that in Period 2 the agent does not know yet his treatment status X. Without this assumption the analysis would be irrelevant, because there would never be any incentive to acquire information on the policy intensity. As we shall see, the analysis can allow for additional time periods and for agents to modify their strategy upon learning their value of X, as long as their behavior before learning X has an effect on the outcome. The effort a involves costs c(a) to be paid in Period 2.

In Period 3, the agent's actual treatment status X is realized. Both a and X affect the outcome Y, which is also realized in Period 3, simultaneously with or after the realization of X. We express the outcome Y given X = x and given a as $W \cdot f(x, a) + \varepsilon$, with $0 \leq f(x, a)$ and $0 < W < \infty$. Here W is just a multiplicative constant in the outcomes, and we merely introduce it to facilitate the analysis of effects of multiplicative changes in the outcomes. In the first leading example, f may be the probability of making an income gain, and W may be the expected income gain. The term ε is an idiosyncratic outcome component. We assume that $\mathbb{E}(\varepsilon|Z) = 0$, but all other determinants of X may be correlated to ε . The latter dependence captures the endogeneity of the actual treatment.

Summarizing, the sequence of events is as follows: the treatment assignment intensity Z is realized, Z is acquired or not, the effort a is chosen, the treatment X is realized, and the outcome $Wf(x, a) + \varepsilon$ is realized.

The above framework gives rise to a "reduced form" treatment evaluation model. First, as will become clear below, we may assume without loss of gener-

⁴The model can be expressed in terms of counterfactual outcomes; see Angrist, Imbens and Rubin (1996) for an exposition.

ality that $\mathbb{E}(X|Z=z)=z$. This leads to a "treatment equation",

$$X = Z + \omega \tag{1}$$

where $\mathbb{E}(\omega|Z) = 0$. Secondly, we have an "outcome equation",

$$Y = Wf(X, a) + \varepsilon \tag{2}$$

with $\mathbb{E}(\varepsilon|Z) = 0$. The analyst observes Y, X, and Z.

Suppose that, more general than equation (1), we would specify $\mathbb{E}(X|Z = z) = g(z)$ for some function g. Typically, it is not difficult to estimate g, and this is why we simply redefine g(z) as our z. This presupposes that g varies with z. In other words, Z as a candidate instrumental variable must be informative. Also, note that the specification $X = Z + \omega$ with $\mathbb{E}(\omega|Z) = 0$ can also capture discrete X. Notably, if X is binary, one may define $X = 1 \iff X^* > 0$ with $X^* = Z + \omega^*$ and ω^* being uniformly distributed on the interval [-1, 0]. To facilitate the exposition we assume that $\Pr(Z \ge 0, X \ge 0) = 1$.

Suppose that in Period 1 the agent does not acquire his personal realization of Z. Then Z only affects Y by way of X, so Z is a valid instrumental variable (IV) because the corresponding exclusion restriction (ER) is satisfied. Now suppose that in Period 1 the agent does acquire his personal realization of Z. Then Z may affect his value of a. In that case, from equation (2), there is a causal effect of Z on the outcome, resulting in a violation of the ER needed for instrumental variable estimation (IVE). Before we analyze this, we first derive in the next subsection the agent's optimal behavior concerning a and concerning acquisition of Z.

3 Economic behavior

3.1 Present values

An agent maximizes his expected present value. To focus on the main issue we consider risk neutral agents. If the agent does not know his value of Z then the expected present value R_0 at the onset of Period 2 equals

$$R_0 = \max_{a \in \mathcal{A}} - c(a) + \frac{1}{1+r} \mathbb{E}_Z \mathbb{E}_{X|Z} \left(W \cdot f(X, a) \right)$$
(3)

 \mathcal{A} denotes the choice set of a, with $\mathcal{A} \subset [0, \infty)$. We denote the optimal a by a_0 .

Now suppose that the agent knows that his value of Z is z. The expected present value R(z) at the onset of Period 2 is

$$R(z) = \max_{a \in \mathcal{A}} -c(a) + \frac{1}{1+r} \mathbb{E}_{X|Z=z} \left(W \cdot f(X,a) \right) \qquad (i = 1, 2)$$
(4)

where r is the discount factor. The optimal a can be expressed as a(z).

The value of information V in Period 1 equals

$$V = \frac{1}{1+r} (\mathbb{E}_Z R(Z) - R_0)$$
 (5)

It is optimal to acquire Z in Period 1 if and only if $V > \gamma$. A central issue of the paper is under which conditions this occurs. For the moment, we simplify the above expressions by subsuming the parameter W/(1+r) into the function f, and 1 + r in (5) into V.

The first insight is that if f is additive in X and a then the optimal a(z) in (4) does not depend on z, and it is equal to a_0 . Consequently, V = 0, and the agent does not acquire Z. In sum,

Proposition 1. If the outcome is additive in the treatment status and the effort of the agent then the exclusion restriction is satisfied.

Throughout the paper we consider functions f that are positive and increasing. As will become clear below, the main practical distinction in the derivations will be whether \mathcal{A} is discrete or not. In the continuous case we often assume that c is quadratic, with $c(a) = \frac{1}{2}c_0a^2$ and $c_0 > 0$. Also, many results will be derived for the following functional form for f,

$$f(x,a) = \psi_0 + \psi_1 a + \psi_2 x + \rho x a$$
(6)

with suitable restrictions on $\psi_1, \psi_2, \psi_3, \rho$ guaranteeing that f is positive and increasing in the relevant intervals for x and a. We do however also provide results for non-parametric specifications of f. The functional form in (6) is concise and allows for explicit expressions and results for the quantities of interest. The interaction parameter ρ captures the degree of complementarity ($\rho > 0$) or substitutability ($\rho < 0$) of treatment and effort, in the outcome. The functional form is less restrictive as may seem. First of all, with binary a and x, (6) is non-parametric because the four possible values of f (i.e., f(1,1), f(1,0), f(0,1), and f(0,0)) are represented by ψ_0, ψ_1, ψ_2 , and ρ . Secondly, as will become clear below, we may generalize the term $\psi_0 + \psi_2 x$ at no cost to a general function $k_2(x)$ (provided that the resulting f is positive and increasing). So all results based on (6) generalize in this respect. The same applies if we replace $\psi_1 a$ by a function $k_1(a)$, provided that $k_1(a) - c(a)$ is quadratic (and again f is positive and increasing), where quadraticness is merely needed to ensure explicit expressions

for the optimal a. More in general, the right-hand side of (6) can be interpreted as the first part of an expansion of the underlying f. It is also useful to point out that in the related literature on decision making with a noisy signal about the unknown state of the world, the general effect of the shape of the outcome function (f) on the value of information (V) is typically too hard to analyze in terms of the model primitives, if no parametric assumptions are made on f (see Persico, 2000, and Athey and Levin, 2001). We return to this literature below.

3.2 Optimal behavior with continuous effort

Suppose that the choice set \mathcal{A} of effort *a* is an interval. We do not restrict *Z* or *X* to be discrete or continuous, so the results below are valid in both cases.

We start with the model in which the functions f and c satisfy (6) and the quadratic specification $c(a) = \frac{1}{2}c_0a^2$, respectively. Within this framework we first consider the case $\rho > 0$. This means that treatment and effort are complements in the outcome. The requirements that f is positive and increasing are then fulfilled by way of the restrictions that $\psi_0 > 0, \psi_1 \ge 0, \psi_2 \ge 0$. Also, we assume that $c_0 > 0$. The optimal a is always positive, and we assume that the lower and upper boundary of \mathcal{A} are not binding for the optimal a.

Let $\overline{z} := \mathbb{E}(Z)$ denote the population mean of Z. It is easy to derive that the optimal effort equals

$$a(z) = \frac{\psi_1 + \rho z}{c_0}$$

$$a_0 = \frac{\psi_1 + \rho \overline{z}}{c_0}$$
(7)

Note that a(z) increases in z. This was to be expected. The complementarity of f implies that the marginal return of effort is higher if the expected (beneficial) treatment level is higher. Also, the optimal effort is higher if the cost of it is lower, if the marginal return of it is higher, and if the degree of complementarity is higher. The optimal a does not depend on the marginal effect ψ_2 of the treatment X.

By substituting (7) into equations (4) and (3), we obtain that

$$V = \frac{1}{2c_0} \rho^2 \operatorname{var}(Z) \tag{8}$$

Consequently, the ER is violated iff $\frac{1}{2c_0}\rho^2 \operatorname{var}(Z) > \gamma$.

Before we interpret this result we first analyze the case $\rho < 0$. This means that treatment and effort are substitutes in the outcome. In the case of participation

in active labor market programs, this case may be more realistic than the case $\rho > 0$. For example, there may be an upper bound on the outcome, and the effort has to compete with efforts for other activities outside of the model.⁵ We also assume again that $c_0 > 0$. The requirement that f is positive and increasing now leads to more complex restrictions on the parameters of f. Notably, the parameters need to ensure that X and a are bounded from above. For a start, we take $\psi_0 > 0, \psi_1 > 0, \psi_2 > 0$. Next, we ensure that f increases in a for every x in the support of X. Sufficient for this is that $\Pr(X < \frac{\psi_1}{-\rho} | Z = z) = 1$ for every z in the support of Z, because this implies that $\Pr(X < \frac{\psi_1}{-a}) = 1$, which implies the desired property of f. Note that $\Pr(X < \frac{\psi_1}{-\rho} | Z = z) = 1$ together with $\mathbb{E}(X|Z=z)=z$ also implies that $\Pr(Z<\frac{\psi_1}{-\rho})=1$, which in turn implies that the optimal a is positive. Finally, we ensure that f increases in X for every a. This is satisfied if $a < \frac{\psi_2}{-\rho}$. We want to have a sufficient condition for this in terms of the model parameters. The optimal a decreases in z, so the largest possible value of a as a function of z is achieved at z = 0. This value equals ψ_1/c_0 . Thus, f has the desired property if $\psi_1/c_0 < \psi_2/(-\rho)$ or, equivalently, $\psi_2 c_0 + \rho \psi_1 > 0$. In sum, we require

$$\psi_0 > 0, \quad \psi_1 > 0, \quad \psi_2 > 0, \quad \psi_2 c_0 + \rho \psi_1 > 0, \quad \Pr(X < \frac{\psi_1}{-\rho} | Z = z) = 1, \quad (9)$$

where the last requirement applies for all z in the support of Z. The requirements also imply that the optimal a satisfies $0 < s \leq \psi_1/c_0$.

We again assume that the lower and upper boundary of \mathcal{A} are not binding for the optimal a. It is not difficult to see that the expressions for the latter are the same as (7), with a(z) now decreasing in z. The resulting expression for V is also the same as in (8). We thus obtain,

Proposition 2. Consider the model with continuous effort, quadratic costs of effort, and the outcome function (6) with the conditions that ensure that it increases in effort and the treatment status. Then the exclusion restriction is violated iff $\rho^2 \operatorname{var}(Z)/(2c_0) > \gamma$.

This means, first of all, that the ER is likely to be violated if the treatment status and the effort before the treatment strongly interact in their effect on the outcome. This is because in that case the optimal effort a(z) is very responsive

⁵This does not rule out that the treatment and the efforts *after* the realization of the treatment are complements in their effect on outcomes in subsequent time periods. These are not in the present model but one can extend it to include them.

to the agent's value z, and the loss of choosing the wrong⁶ amount of effort is larger. For example, if $\rho > 0$ then knowing that z is large leads to an optimal effort a(z) that is also very large, while knowing that z is small leads to a small a(z), and in both of these cases the alternative choice of an intermediate effort level a_0 entails a substantial loss.

Violation is also more likely if Z has a large variance. In that case, the candidate instrument generates a large range of mean policy intensities X|Z, and it is more likely that not acquiring Z leads to a large loss. If the effort cost parameter c_0 is large then violation is less likely, because then the optimal effort is small whether one acquires z or not. Violation of the ER is more likely if γ is small, which is trivial to understand.

It is also useful to discuss which model parameters do *not* affect the likelihood that the ER is violated. First, consider the parameter ψ_2 capturing the additive treatment-effect component that does not depend on effort. It does not affect the value of information precisely because the additive effect cannot be influenced by the effort. Until now we have not defined summary treatment effects yet. However, it is clear that any such measure depends on ψ_2 . We conclude that the size of the (average) treatment effect does not affect the validity of the ER.

Secondly, consider the strength of the candidate instrument. This is usually defined as the strength of the association between Z and X, for example as measured by the squared correlation coefficient $R^2(X,Z)$, which in our model reduces to $\operatorname{var}(Z)/\operatorname{var}(X)$. This quantity does not have a direct effect on the validity of the ER. If the residual variance $\operatorname{var}(\omega)$ in the "treatment equation" (1) is of a higher order of magnitude than $\operatorname{var}(Z)$ then the candidate instrument is weak but it may nevertheless have a large variance by itself, and the latter makes it likely that the ER is violated. The underlying reason is that agents are only concerned with the *mean* of the treatment status X given z, when they decide on their optimal effort a.⁷

$$V = \frac{1}{2c_0} \rho^2 \operatorname{var}_Z(\mathbb{E}(k(X)|Z))$$

If, for example, $k(x) = x^3$ and $\omega \equiv X - Z$ is symmetric, then $\operatorname{var}_Z(\mathbb{E}(k(X)|Z))$ can be shown to equal $\operatorname{var}(Z^3) + 9(\operatorname{var}(\omega))^2 \operatorname{var}(Z)$, which in turn equals

$$\operatorname{var}(Z^3) + 9(\operatorname{var}(Z))^3 \left[\frac{1}{R^2(X,Z)} - 1\right]^2$$

 $^{^{6}\}mathrm{In}$ this section, "wrong" is used in the sense of "wrong if the agent had known his actual value of Z ".

⁷More general functional forms of f(x, a) may have somewhat different implications. If $f(x, a) = \rho k(x)a$ for some increasing function k then it can be shown that

Thirdly, the validity of the ER does not depend on ψ_0 and ψ_1 . The fact that it does not depend on any of the ψ_i parameters reflects the fact that additive effects of treatment status and effort do not affect the ER (recall Proposition 1).

We now return to the related literature on decision making with a noisy signal about the state of the world. In this literature, agents receive a signal (say, Z) about the unknown state of the world (say, X) and they have to decide on which action (say a) to take. The outcome (say, f) depends on X and a (see Gollier, 2001, for a recent overview of models with signals, effort, and outcomes). The main difference with our setup is that in this literature the focus is on the strength of the causal effect from the state of the world X on the signal Z (or, equivalently, on the quality of the signal), whereas in our setup Z causally affects X. Also, this literature restricts attention to outcome functions f that satisfy generalized notions of complementarity in the state of the world and the action of the individual.

Nevertheless, some of the results from this literature are directly applicable to our context. Athey and Levin (2001) present a generalized version of the following. Consider our model with continuous effort, where the treatment status X increases in Z in the sense of first-order stochastic dominance. Let the outcome function increase in effort and in the treatment status, and let the cross-derivative be positive. Then there is a monotone positive relation between z and the optimal effort a(z). If no parametric assumptions are made on f then the general effect of the shape of the outcome function (f) on the value of information (V) is typically too hard to analyze in terms of the model primitives (see Persico, 2000, and Athey and Levin, 2001). Results that have been derived for given (i.e., not optimally determined) functions a(z) emphasize the importance of the degree of complementarity of the outcome function for the value of information. We return to this in Subsection 4.2.

3.3 Optimal behavior with discrete effort

Now let \mathcal{A} be discrete. For expositional reasons, we simplify the analysis by assuming that a is binary, i.e. is taken from the set $\{0, 1\}$. We adopt specification (6) for f. The cost of effort function c(a) is now represented by its two possible values c(0) and c(1). By analogy to the previous subsection, we denote c(1) - c(0) by c_0 , which may be called *the* cost of effort. The expressions for the present values can now be simplified to

Therefore, the value of information V increases in $var(\omega)$, and increases in the weakness of the candidate IV as conventionally defined.

$$R(z) = \psi_0 - c(0) + \psi_2 z + \max\{0, \psi_1 + \rho z - c_0\}$$
(10)

$$R_0 = R(\overline{z})$$

and the agent chooses a = 1 iff the second term in the maximum exceeds 0, so we replace equation (7) by

$$a(z) = I\left(\frac{\psi_1 + \rho z}{c_0} > 1\right)$$

$$a_0 = I\left(\frac{\psi_1 + \rho \overline{z}}{c_0} > 1\right)$$
(11)

where I(.) denotes the indicator function. The intuition behind these expressions is exactly as for equation (7). The difference is that a(z) is now discrete instead of continuous, and the optimal choice of a is now insensitive to small changes in the parameter values.

To facilitate the exposition, we make the additional assumption that Z has two possible values in the population of agents: $\Pr(Z = z_1) = p = 1 - \Pr(Z = z_2)$, with $z_1 \neq z_2$ and $0 and normalization <math>z_1 > z_2$. It is again useful to distinguish between the two cases in which treatment and effort are substitutes $(\rho < 0)$ or complements $(\rho > 0)$ for the outcome. This time we start with the former case, which requires again some restrictions on the range of values of the model parameters. By analogy to the previous subsection, we impose

$$\psi_0 > 0, \quad \psi_1 > 0, \quad \psi_2 > 0, \quad \psi_2 + \rho > 0, \quad \Pr(X < \frac{\psi_1}{-\rho} | Z = z_i) = 1$$

where the last requirement applies for both z_i . In fact, the results below also apply if $\psi_2 + \rho = 0$, meaning that with effort a = 1 it is irrelevant whether the treatment is realized or not. This describes situations in which the outcome of interest is the transition from unemployment to work if this is achieved with certainty by the effort a = 1 but the treatment by itself cannot achieve this.

With $\rho < 0$, equation (11) implies that

$$0 \le a(z_1) \le a_0 \le a(z_2) \le 1$$

and we can distinguish between the following four "regimes",

• Regime 1. $\psi_1 + \rho z_1 > c_0$. Then $a(z_1) = a_0 = a(z_2) = 1$.

- Regime 2. $\psi_1 + \rho z_1 \le c_0 < \psi_1 + \rho \overline{z}$. Then $0 = a(z_1) < a_0 = a(z_2) = 1$.
- Regime 3. $\psi_1 + \rho \overline{z} \le c_0 < \psi_1 + \rho z_2$. Then $0 = a(z_1) = a_0 < a(z_2) = 1$.
- Regime 4. $\psi_1 + \rho z_2 \leq c_0$. Then $0 = a(z_1) = a_0 = a(z_2)$.

Next, we consider V. In Regime 1, V = 0, implying that the information is not bought. This is because the agent knows that he will always provide effort, under every policy, whether the policy is known or not, so the information is irrelevant for the optimal behavior. The same line of reasoning applies to Regime 4. So, for certain extreme parameter values, the agent does not acquire his value of Z. Now consider Regimes 2 and 3. We obtain that in Regime 2,

$$V = p\left(c_0 - \psi_1 - \rho z_1\right)$$

(recall that $p := \Pr(Z = z_1)$). This result is particularly easy to interpret. In Regime 2, the information on Z is valuable if and only if in truth the agent has policy intensity z_1 , because he *does* provide effort if he has no information. Therefore the value V equals minus the expected loss of making such a wrong decision. More precisely, V equals the product of [the probability that, when Z is not acquired, an effort a_0 is chosen that is not optimal in the light of the actual Z] and [minus the loss of choosing a_0 given that the actual Z would lead to another choice of a]. In Regime 2, the statement in the first square brackets equals $\Pr(Z = z_1)$, which is p. The statement in the second set of square brackets equals [the additional cost c_0 of choosing $a_0 = 1$ compared to the costs made if one had known that $Z = z_1$] minus [the additional expected return of choosing $a_0 = 1$ compared to the expected return if one had known that $Z = z_1$]. Similarly, in Regime 3,

$$V = (1 - p) \left(\psi_1 + \rho z_2 - c_0\right)$$

V > 0 if and only if the optimal effort when knowing that $Z = z_1$ differs from the optimal effort when knowing that $Z = z_2$. This is intuitively clear: only in these cases does the knowledge of Z potentially have an effect on the effort provided.

The results for the case $\rho > 0$ are a mirror image of those for $\rho < 0$. The optimal efforts satisfy $0 \le a(z_2) \le a(\overline{z}) \le a(z_1) \le 1$, and these can be characterized in terms of the model parameters by applying equation (11). For sake of brevity we do not present the other results for this case.

Consider Figure 1, plotting V against c_0 for the case $\rho < 0$. The maximum of V as a function of c_0 is attained at the boundary between Regimes 2 and 3, which means $c_0 = \psi_1 + \rho \overline{z}$. This value V_{max} equals



Figure 1. The value of information *V* as a function (indicated by the fat solid curve) of the cost of effort $c_0 := c(1)-c(0)$.

Note: the parameter γ denotes the information acquisition cost.

$$V_{max} = |\rho|p(1-p)(z_1 - z_2)$$

which also holds if $\rho \geq 0$. $V_{max} > \gamma$ iff there are values c(1), c(0) of the cost-ofeffort function for which it is optimal to buy the information. The expression for V_{max} is similar to the expression (8) for V in the continuous case. In particular, note that $\operatorname{var}(Z) = p(1-p)(z_1-z_2)^2$. The comparative statics results concerning the effects of ρ, γ and $\operatorname{var}(Z)$ for the continuous case therefore carry over to the present discrete case. For example, concerning p we find that acquisition is more likely if p is not too large and not too small, because p = 1/2 maximizes the variance of Z in the population, for given values of z_1, z_2 . The a priori uncertainty concerning which value of Z applies is largest for such intermediate values of p, and so is therefore the a priori probability of providing the wrong amount of effort.

What is particular about the discrete case is that for certain parameter values (namely in Regimes 1 and 4) the optimal a does not depend on Z. Due to the continuity of the value functions, the ensuing low value of information also applies for parameter values such that the agent is in one of the other regimes but close to Regimes 1 and 4. In those cases the loss of making a wrong decision on a is too small to justify the acquisition of one's value of Z. Consequently, it is important

to assess which parameter values lead to Regimes 1 or 4. Notably, the agent is more likely to be in one of the latter regimes if c_0 is very large or very small. The underlying explanation is that for such high or low values of c_0 , acquiring information on Z would never lead to behavioral changes. Similarly, if ψ_1 is very large while the other parameters have moderate values then effort has such a positive effect on the outcome that the agent always provides effort a = 1.

We may re-interpret the effects of the costs parameters c_0 and γ in terms of marginal utility losses for a risk-averse agent who cannot transfer resources between time periods and who has a per-period utility function that displays decreasing absolute risk aversion. If the agent has a low per-period baseline income then the marginal utility losses of costs c_0 in Period 2 and γ in Period 1 are relatively high. The acquisition of information is unlikely for two reasons: it is expensive and it is useless because it would not affect optimal effort. This makes violation of the exclusion restriction unlikely. If the agent has a very high per-period baseline income then the marginal utility loss of costs c_0 in Period 2 is small, and acquisition is again unlikely, but now only for the second of the above two reasons.⁸ Again, the exclusion restriction is then relatively easy to justify. For agents with per-period baseline incomes in between these extremes it is more likely that information on Z is acquired, so then the exclusion restriction is more easily violated. From this point of view, instrumental variable analyses that restrict attention to agents with low or high resources (e.g. income) are more likely to be valid than analyses that include agents with intermediate resource levels, if effort is binary.

One may combine the discrete case of this subsection with the continuous case of the previous subsection, e.g. by allowing the effort a to attain all values in a fixed interval, e.g. [0, 1], with bounds that are binding for certain model parameter values. Sufficiently large parameter value changes then typically lead to results that correspond to those in this subsection. An alternative theoretical extension would be to allow agents to manipulate the probability distribution of their treatment status X by way of their choice of a. In Appendix 1 we show how this can be translated into our model framework.

We now return to the leading example, where the treatment is participation in a job search assistance program for unemployed individuals, and the policy intensity differs across regions. An unemployed individual knows that there is a

⁸Unless of course the expected marginal return of providing effort is always very small as well. The problem with characterizing comparative statics effects on optimal behavior if the latter is a highly non-linear function of a large number of parameters is that there are often joint limiting values of a subset of parameters that "push" the result in any desired direction.

chance that he may enroll in a job search assistance program, and this may affect his job search strategy before actual enrollment. If the program is attractive and if the applicable regional probability of enrollment into the program is known, then this probability may affect the optimal private job search effort before enrollment into the program. With a negative interaction effect of treatment and effort on the outcome of being employed within a year, the effort is lower if the probability of enrollment is higher. We now apply the results in order to inquire the conditions under which Z is a valid instrumental variable, that is, the conditions under which the agent does not acquire his value of Z. First, it may be relatively easy for agents to learn the specific situation in their own region and the effect of this on the rate at which they expect to be treated. This should lead to caution when using regional variation in the budget for (or, more generally, an indicator of the geographical availability of) active labor market programs as an instrumental variable to study causal effects of the program. The agents' incentive to acquire information is particularly high if treatment and effort are close substitutes in the outcome, or if the policy intensity varies strongly across the regions, or if the agents take the a priori probability of whether one lives in the high-budget region to be equal to around 1/2 which means that about half of the agents are exposed to the high-budget situation.

Note that the case where Z and Y are binary is equivalent to a randomized experiment with deliberate randomization. Typically, in experiments, the intention to treat is randomized with probability equal to 1/2. Our results suggest that (provided that the individual randomization outcome can only be acquired at a positive cost when one has to decide the effort level a) it may be better to use a smaller or larger probability, because this reduces the likelihood that the agent has an incentive to acquire and use the value of the candidate instrumental variable.

So far we have only examined a single candidate instrumental variable. Clearly, if there are many such variables, each giving only limited information on the treatment assignment process, and if the cost of information acquisition is linear in the number of variables on which information is acquired, then this reduces the likelihood that the exclusion restriction is violated. Other obvious results concern the timing of events (i.e., the relative lengths of the time periods in the model). For example, with a small amount of time between the moment at which the policy intensity is determined and the moment at which the treatment is realized, the scope for information acquisition is reduced. If treatment and outcome are realized close in time then acquisition is attractive from the point of view that its expected future returns are discounted less heavily.

4 The magnitude of the bias of the instrumental variable estimator if the exclusion restriction is violated

4.1 The parameter of interest and the estimator

In this section we address the bias of the instrumental variable estimator due to the exclusion restriction violations that we consider. It is useful to return to the "reduced form" model representation from Section 2. IVE involves the estimation of the effect of X on Y, holding all other determinants of Y constant. In the model, this is the partial derivative or first difference of f with respect to its first argument. The classical IV regression estimator, if applied to data on Y, X, and Z, estimates $\operatorname{cov}(Y, Z)/\operatorname{cov}(X, Z)$. More precisely, estimation involves that these two covariances are replaced by their sample equivalents, and then the probability limit of the estimator, which we denote by $\widehat{\beta}_{IV}$, satisfies

$$\widehat{\beta}_{IV} \to \frac{\operatorname{cov}(Y, Z)}{\operatorname{cov}(X, Z)}$$

In our model framework, if the ER is valid, this equals

$$\frac{\operatorname{cov}(f(X, a_0), Z)}{\operatorname{cov}(X, Z)} \tag{12}$$

We simply define this to be the parameter of interest β . It is a summary measure of the slope of the average outcome as a function of treatment status for a given fixed effort. The focus on β is particularly sensible if f is linear in X, as is the case in specification (6), because then this slope is constant and equal to β .⁹

⁹For more general f one may consider alternative parameters of interest. First, f as a function of x may be estimated non-parametrically. If X and Z are continuous then f can be approximated by a power series and estimated with series estimation using instruments $Z, Z^2, Z^3...$ (see Blundell and Powell, 2003, for an overview). Secondly, one may focus on the average derivative $AD := \mathbb{E}(\frac{\partial f}{\partial x}(X, a_0))$ as parameter of interest. In both cases the estimates are biased if the actual effort a depends on Z. However, formal insightful analyses of the bias for general f are hampered by complexity of the expressions involved.

In fact, under certain conditions, $\beta = AD$. Notably, if a general function $f(x, a_0)$ as a function of x can be represented by a power series, then $\beta = AD$ if (a) the second and higher order coefficients are all zero, or if (b) the third and higher order coefficients are all zero and Z is symmetrically distributed, or if (c) Z is normally distributed (which however is formally ruled out by $\Pr(Z \ge 0) = 1$). For general f, application of Stein's Lemma provides a justification for β as an approximation of the average derivative. Notably, for two random variables U_1, U_2 and a function g, this Lemma states that $\operatorname{cov}(g(U_1), U_2) \approx \mathbb{E}(g'(U_1))\operatorname{cov}(U_1, U_2)$, so application

Moreover, as we shall see, β is a local average treatment effect if Z is discrete with two points of support. Note that we assume population homogeneity of all model parameters, functions, and distributions.¹⁰

Throughout the upcoming subsections we assume that the function f governing the observed outcomes is the same as the function f that agents use in Periods 1 and 2 to assess their expected utility in Period 3. This entails that agents know the average treatment effect that the researcher aims to estimate. In Subsection 4.5 we relax this assumption and show how the results change.

Expression (12) can be simplified by noting that cov(X, Z) equals var(Z). Also, for any two random variables U_1, U_2 , there holds that $cov(U_1, U_2) = cov(\mathbb{E}(U_1|U_2), U_2)$. This results in

$$\beta = \frac{\operatorname{cov}(\mathbb{E}(f(X, a_0)|Z), Z)}{\operatorname{var}(Z)}$$
(13)

If the ER does not apply, then

$$\widehat{\beta}_{IV} \to \frac{\operatorname{cov}(f(X, a(Z)), Z)}{\operatorname{cov}(X, Z)} = \frac{\operatorname{cov}(\mathbb{E}(f(X, a(Z))|Z), Z)}{\operatorname{var}(Z)}$$
(14)

Note that $\widehat{\beta}_{IV}$ captures the over-all effect of Z on the outcome. If the ER is violated then the over-all effect does not equal the causal treatment effect but also includes the causal chain that runs by way of the effort a. From equations (13) and (14) it follows that

$$\widehat{\beta}_{IV} - \beta \to \frac{\operatorname{cov}(\mathbb{E}(f(X, a(Z)) - f(X, a_0) | Z), Z)}{\operatorname{var}(Z)}$$
(15)

Somewhat loosely, this is an average of the effect on the outcome of the dependence of a on Z.¹¹ If treatment and effort are complements in the outcome then one may expect that the optimal effort is an increasing function of z, and

leads to $\beta \approx \mathbb{E}(\frac{\partial f}{\partial x}(X, a_0)).$

¹⁰Any meaningfully defined treatment effect will depend on effort a_0 which will in general depend on properties of the distribution of Z like $\mathbb{E}(Z)$ (note that this concerns situations in which the ER is satisfied). In empirical studies, such an association between the treatment effect and the degree at which a population is exposed to a policy can be mistakenly attributed to equilibrium effects of the policy.

¹¹This is again more transparent for approximations of the asymptotic bias. Application of Stein's Lemma (see footnote 9) leads to $\hat{\beta}_{IV} - \beta \approx \mathbb{E}_Z(\frac{d}{dz}\mathbb{E}(f(X, a(Z)) - f(X, a_0)|Z))$). Application of the Delta method approximation (which, for a random variable U and a function g, states that $\operatorname{cov}(g(U), U) \approx g'(\mathbb{E}(U))\operatorname{var}(U)$) leads to $\hat{\beta}_{IV} - \beta \approx \frac{d}{dz}\mathbb{E}(f(X, a(z)) - f(X, a_0)|z)|_{z=\overline{z}}$. It can be shown that if z is a location parameter of the distribution of [X|Z=z] and if $a(\overline{z}) = a_0$ and a(z) is differentiable then this simplifies further to

consequently one may expect the asymptotic bias to be positive. In the next subsections we examine this more formally.

If Z describes the assigned treatment (as opposed to the actual treatment X) then the over-all effect is usually called the intention-to-treat (ITT) effect on the outcome. This can be decomposed into the actual treatment effect β and the announcement or ex ante effect of the treatment (see e.g. Abbring and Van den Berg, 2003, 2005, for this terminology). The latter thus equals $\hat{\beta}_{IV} - \beta$, which is the asymptotic bias of the IV estimator $\hat{\beta}_{IV}$ of β .

From an econometric regression point of view, one may state that the asymptotic bias of the IV estimator results from the fact that the size of the causal effect of treatment on outcome depends on the candidate IV. An alternative way to look at the asymptotic bias is to write the outcome equation as Y = $f(X, a_0) + (f(X, a(Z)) - f(X, a_0)) + \varepsilon$. By ignoring the dependence of a on Z, an IV regression analysis takes the sum of the second and third terms in the right-hand side as the residual term in the outcome equation. Consequently, the candidate IV Z is correlated to the error term in the outcome equation.

Note that instrumental variable estimators are typically biased in case of ER violations even if the residual terms ε and ω in the outcome and treatment equations are independent, i.e. even if the treatment is not selective in the usual sense of the word.

4.2 Continuous effort

Suppose that the ER is violated. We are first going to examine the asymptotic bias in the model with continuous effort, quadratic costs of effort, and the outcome function (6), with the conditions that ensure that it increases in effort and the treatment status. Subsequently we generalize the cost-of-effort function and the outcome function.

With the outcome function (6), β as defined above is the average treatment effect in the population $\partial f(x, a)/\partial x$ which in this specific case does not depend on x, and which is evaluated at $a = a_0$. By substituting f into (13), and substituting a_0 , we obtain,

$$\widehat{\beta}_{IV} - \beta \approx a'(\overline{z}) \mathbb{E}_{\omega} \left(\frac{\partial f}{\partial a} (\omega + \overline{z}, a_0) \right)$$

where $\omega = X - Z$ (see equation (1)). Clearly, $a'(\overline{z})$ captures the responsiveness of effort to the value of Z, and the second term in the right hand side captures the effect of effort on the outcome. We return in the next subsection to the quality of these approximations in special cases.

$$\beta = \psi_2 + \rho a_0 = \psi_2 + \rho \frac{\psi_1 + \rho \overline{z}}{c_0}$$
(16)

which is always positive, by virtue of the conditions ensuring that f is increasing.

Similarly, by substituting (6) and a(z) from (7) into (14), we obtain,

$$\widehat{\beta}_{IV} \to \psi_2 + \frac{2\rho\psi_1}{c_0} + \frac{\rho^2}{c_0} \frac{\operatorname{cov}(Z, Z^2)}{\operatorname{var}(Z)}$$
(17)

As a result,

$$\widehat{\beta}_{IV} - \beta \to \frac{\rho}{c_0} \left[\psi_1 - \rho \overline{z} + \rho \frac{\operatorname{cov}(Z, Z^2)}{\operatorname{var}(Z)} \right]$$
(18)

In Appendix 2 we prove the following:

Proposition 3. Consider the model with continuous effort, quadratic costs of effort, and the outcome function (6) with the conditions that ensure that it increases in effort and the treatment status. If the exclusion restriction is violated then the asymptotic bias of the IV estimator $\hat{\beta}_{IV}$ has the same sign as ρ .

This result reinforces earlier results on the importance of the interaction parameter. For $\rho > 0$ the result is very intuitive. In that case, if the agent acquires his value z of Z, then effort is increasing in z; the estimated treatment effect is then boosted, and the causal effect of the treatment is over-estimated, so asymptotically $\hat{\beta}_{IV} > \beta$. In Appendix 2 we demonstrate that if $\rho > 0$ then $\hat{\beta}_{IV}$ and $\hat{\beta}_{IV} - \beta$ increase in ρ . If ρ increases then it is optimal to make the effort a(z) more responsive to z, so that the announcement effect and therefore the asymptotic bias increase. For $\rho < 0$ the additive and the interaction effects of a(z) have opposite signs, but the former dominates. Note also that the asymptotic bias decreases in c_0 . If costs of effort are low then it is pays off to supply a large amount of effort for certain z, thus making effort more responsive to z.

In Appendix 2 we also demonstrate that (18) can be rewritten as

$$\widehat{\beta}_{IV} \to 2\beta - \psi_2 + \frac{\rho^2 \overline{z}}{c_0} \xi \sigma \tag{19}$$

where ξ and σ are the skewness and the coefficient of variation of Z, so $\xi := \kappa_3/\kappa_2^{\frac{1}{2}}$ and $\sigma := \kappa_2^{\frac{1}{2}}/\overline{z}$, with $\kappa_i := \mathbb{E}(Z - \overline{z})^i$. From this equation some additional comparative statics results follow. As β does not depend on moments of Z higher than the first, it follows that $\hat{\beta}_{IV}$ and $\hat{\beta}_{IV} - \beta$ increase in the skewness ξ of Z. This makes sense: e.g. if Z is highly skewed to the right then relatively many agents have a very high treatment value and a very high effort level, leading to a very high correlation between outcome and candidate instrument.

Notice that if Z is symmetric (i.e. $\xi = 0$) then the variance of Z does not affect the asymptotic bias of $\hat{\beta}_{IV}$. In this case, increasing $\operatorname{var}(Z)$ leads to an increasing value of information, and therefore an increasing likelihood of ER violation, but not to an increasing asymptotic bias of the IV estimator.¹² So in this sense the asymptotic bias is even less sensitive to the strength of the candidate instrument than the ER.

Equation (19) can also be used to bound β further. We consider using the sign of the skewness of Z, which is always observable. In particular, in case of a balanced experiment, $\xi = 0$. We also consider cases where one has a priori knowledge on the sign of ρ and on whether the additive effect ψ_2 of treatment on outcome is zero or positive.

Corollary 1. Under the conditions of Proposition 3,

- (i). If Z is symmetric or skewed to the left then, asymptotically, $\beta \geq \frac{1}{2}\widehat{\beta}_{IV}$.
- (ii). If Z is symmetric and $\psi_2 = 0$ then, asymptotically, $\beta = \frac{1}{2}\widehat{\beta}_{IV}$.
- (iii). If Z is skewed to the right and $\psi_2 = 0$ then, asymptotically, $\beta < \frac{1}{2}\widehat{\beta}_{IV}$.
- (iv). If Z is symmetric or skewed to the left and $\rho > 0$ then, asymptotically, $\frac{1}{2}\widehat{\beta}_{IV} \leq \beta < \widehat{\beta}_{IV}$.

Extension to other cases is straightforward. Notice that $\psi_2 = 0$ implies that $\rho > 0.^{13}$

We now consider more general models. First, if we replace the additive term $\psi_0 + \psi_2 x$ in the outcome function by a general function $k_2(x)$ then all asymptotic

$$\widehat{\beta}_{OLS} = \widehat{\beta}_{IV} \cdot R^2(X, Z) + \beta \cdot (1 - R^2(X, Z)) + \frac{\operatorname{cov}(\varepsilon, \omega)}{\operatorname{var}(X)}$$

¹²With our outcome and cost functions, the case $\xi = 0$ is also the case for which the asymptotic bias approximations by Stein's Lemma and the Delta method, in footnote 11, are exact.

¹³Comparison of the performance of the IV estimator to that of the OLS estimator is not very informative, as the latter is also driven by the covariance of the residual terms ε and ω whereas the former is not. It can be shown that asymptotically,

which holds regardless of whether the ER is satisfied or not (in the former case, of course, $\hat{\beta}_{IV} = \beta$). It follows that IV does better than OLS if the IV bias due to violation of the ER is small and the "selection effect" as measured by $\operatorname{cov}(\varepsilon, \omega)$ is large. In absence of such a selection effect, OLS does better than IV if the ER is violated and $R^2(X, Z) < 1$. With randomized experiments, $X \equiv Z$ and $\varepsilon \perp \omega$, so $\hat{\beta}_{OLS} = \hat{\beta}_{IV}$.

bias results remain valid. Next, we relax the assumption that costs of effort are quadratic.

Proposition 4. Consider the model with continuous effort and the outcome function (6) with $\psi_0 > 0, \psi_1, \psi_2 \ge 0$, and $\rho > 0$. Let the cost of effort be increasing, convex, and differentiable and lead to an interior solution for the optimal effort. If the exclusion restriction is violated then the asymptotic bias of the IV estimator $\hat{\beta}_{IV}$ is positive.

See Appendix 3 for the proof.

It is more difficult to extend the result to the case where treatment and effort are substitutes in the outcome (i.e., $\rho < 0$). This is because with $\rho < 0$ the additive treatment effect and the interaction effect of treatment and effort have opposite signs.¹⁴

Next, we relax the assumption that f satisfies (6), so we do not make any parametric assumptions on the cost function and the outcome function. We use the monotonicity result of Athey and Levin (2001) (see Subsection 3.2) as an input. As noted above, in the literature on decision making with signals on the unknown state of the world, comparative statics are sometimes too hard to analyze in terms of the model primitives, and one can only derive results for given (i.e., not optimally determined) effort functions. To a certain extent our next result shares this feature, in that we assume that the optimal effort function a(z) satisfies $a(\overline{z}) = a_0$ without translating this into model primitives. The assumption means that the agent's optimal effort if he knows that he has the average treatment intensity is equal to the agent's optimal effort if he does not know his intensity. In the cases considered in the previous propositions, this assumption actually follows from the assumptions on the model primitives.

Proposition 5. Let the outcome function f(x, a) be non-negative, increasing and supermodular in x and a. Let the treatment status X increase in Z in the sense

$$\widehat{\beta}_{IV} - \beta \approx (\psi_1 + \rho \overline{z}) a'(\overline{z})$$

¹⁴See e.g. equation (33) in Appendix 3. Specifically, with less parametric assumptions, it is cumbersome to formulate and exploit restrictions ensuring that the outcome increases in treatment and effort everywhere if treatment and effort are substitutes. However, using the Delta method approximation (see footnote 11) the desired result follows. We substitute the result from Appendix 3 that $a(\bar{z}) = a_0$ into the approximation, to obtain:

With $\rho < 0$, the function a(z) is decreasing (see Appendix 3), and the requirement that f is increasing implies that $\overline{z} < \psi_1/(-\rho)$. The right-hand side of the above approximation is therefore negative.

of first-order stochastic dominance. Let the optimal effort function a(z) satisfy $a(\overline{z}) = a_0$. If the exclusion restriction is violated then the asymptotic bias of the IV estimator $\widehat{\beta}_{IV}$ is non-negative.

See Appendix 4 for the proof. Supermodularity captures complementarity and is satisfied if the cross-derivative of f(x, a) with respect to x and a is positive. The assumption that X increases in Z in the sense of first-order stochastic dominance is not nested with the assumption that $\mathbb{E}(X|Z = z) = z$, although the former is weaker as a condition for $\mathbb{E}(X|Z = z)$ to increase in z.

Note that Proposition 5 does not require continuity of effort a. In the next subsection we consider the discrete case in more detail.

4.3 Discrete effort

Let a be binary and let Z have a discrete distribution with points of support $z_1 > z_2$, like in Subsection 3.3. It can be shown that β in (13) then simplifies to

$$\beta = \frac{\mathbb{E}(f(X, a_0)|z_1) - \mathbb{E}(f(X, a_0)|z_2)}{z_1 - z_2}$$

which is a local average treatment effect (compare Imbens and Angrist, 1994; note that the denominator equals $\mathbb{E}(X|z_1) - \mathbb{E}(X|z_2)$). If the ER is valid then the IV estimator $\hat{\beta}_{IV}$ converges to this number. Accordingly, a Wald estimator can be used as IV estimator.

If the ER is violated then we can simplify (14) to

$$\widehat{\beta}_{IV} \to \frac{\mathbb{E}(f(X, a(z_1))|z_1) - \mathbb{E}(f(X, a(z_2))|z_2)}{z_1 - z_2}$$
(20)

so that again $\widehat{\beta}_{IV}$ captures the over-all effect of Z on the outcome.

We proceed by taking the outcome function f to satisfy (6). From the above (as well as from the previous subsection) it immediately follows that $\beta = \psi_2 + \rho a_0$. If the ER is violated then necessarily $a(z_1) \neq a(z_2)$. With $\rho < 0$, violation implies that $0 = a(z_1) < a(z_2) = 1$, whereas with $\rho > 0$ this is reversed. By elaborating on equation (14) or on equation (20), we obtain

$$\begin{aligned} \widehat{\beta}_{IV} &\to \psi_2 - \frac{\psi_1 + \rho z_2}{z_1 - z_2} & \text{if } \rho < 0\\ \widehat{\beta}_{IV} &\to \psi_2 + \frac{\psi_1 + \rho z_1}{z_1 - z_2} & \text{if } \rho > 0 \end{aligned}$$

Depending on the sign of ρ and the value of a_0 , we have four different expressions for the asymptotic bias,

$$\begin{aligned} \widehat{\beta}_{IV} - \beta &\to -\frac{\psi_1 + \rho z_2}{z_1 - z_2} & \text{if } \rho < 0 \text{ and } a_0 = 0 \\ \widehat{\beta}_{IV} - \beta &\to -\frac{\psi_1 + \rho z_1}{z_1 - z_2} & \text{if } \rho < 0 \text{ and } a_0 = 1 \\ \widehat{\beta}_{IV} - \beta &\to \frac{\psi_1 + \rho z_1}{z_1 - z_2} & \text{if } \rho > 0 \text{ and } a_0 = 0 \\ \widehat{\beta}_{IV} - \beta &\to \frac{\psi_1 + \rho z_2}{z_1 - z_2} & \text{if } \rho > 0 \text{ and } a_0 = 1 \end{aligned}$$

We thus obtain,

Proposition 6. Consider the model with binary effort, a candidate IV with a discrete distribution with two points of support, and the outcome function (6) with the conditions that ensure that it increases in effort and the treatment status. If the exclusion restriction is violated then the asymptotic bias of the IV estimator $\hat{\beta}_{IV}$ has the same sign as ρ .

which is very similar to Proposition 3 of the previous subsection. The intuition is the same.

The bias terms listed above increase in absolute value if $z_1 - z_2$ decreases. However, recall that the expressions are derived under the assumption that the ER is violated. For fixed values of the other parameters, the ER is satisfied for sufficiently small values of $z_1 - z_2$.

4.4 Testing for a causal treatment effect in absence of an exclusion restriction

Despite the fact that the IV estimator is asymptotically biased if the ER is violated, we can still use standard tests to inquire whether the treatment has a causal effect on the outcome. To see this, notice that in absence of a treatment effect the value of information is zero, so there is no acquisition of information, implying that the ER is satisfied and the IV treatment effect estimator is asymptotically equal to zero. The standard asymptotic tests of the null hypothesis of a zero treatment effect therefore have the correct size. This exploits the insight that in our model any violation of the ER is behaviorally triggered by a non-zero treatment effect.

One may wonder what the power is of such tests if the null hypothesis of no treatment effect is incorrect in reality and, in addition, the ER is violated. We shed some light on this by examining the (asymptotic) signs of β , $\hat{\beta}_{IV}$, and $\widehat{\beta}_{IV} - \beta$. According to expression (13), the treatment effect parameter β has the same sign as $\operatorname{cov}(\mathbb{E}(f(X, a_0)|Z), Z)$. If the treatment status X increases in Z in the sense of first-order stochastic dominance, and f(x, a) strictly increases in x, then $\mathbb{E}(f(X, a_0)|Z)$ increases in Z, and consequently $\beta > 0$.

To proceed, it matters again whether treatment and effort are complements or substitutes in the outcome. For the former case we saw earlier in this section that asymptotically $\hat{\beta}_{IV} \geq \beta$. Therefore one may expect a higher power of the standard asymptotic tests of the null hypothesis of a zero treatment effect, compared to when the ER is not violated. This makes sense: the true treatment effect is magnified by the agents' efforts.

For the substitution case with our parametric functional forms of f(x, a) and c(a), we saw that asymptotically $\hat{\beta}_{IV} < \beta$. The next example shows that it is even possible that asymptotically $\hat{\beta}_{IV} = 0$.

Example 1. Let $\psi_1 = \psi_2 = c_0 = 1$ and $\rho = -3/4$. Note that these values satisfy the parameter restriction $\psi_2 c_0 + \rho \psi_1 > 0$ (see (9)). In turn, they imply that $\Pr(Z < 4/3) = 1$. By substituting our f and a(z) into (14) we obtain that asymptotically $\hat{\beta}_{IV} \ge 0$ iff $\operatorname{cov}(Z, -\frac{1}{2}Z + \frac{9}{16}Z^2) \ge 0$. The function $-\frac{1}{2}Z + \frac{9}{16}Z^2$ strictly decreases in Z on the interval $(0, \frac{4}{9})$, so if this interval includes the support of Z then asymptotically $\hat{\beta}_{IV} < 0$. If Z has a discrete distribution with as only two points of support 1/3 and 5/9 then $\hat{\beta}_{IV} \to 0$.

In this pathological case the asymptotic power is zero. In general, in the substitution case, one may expect the power to be lower if the ER is violated than if the ER is satisfied.

One way to improve the power is to look for evidence of higher-order dependencies between the candidate instrument Z and the outcome Y. Under the null hypothesis, these must be absent as well.

4.5 Systematic difference between the outcome and the agent's utility in the outcome period

We now allow the function f^* governing the observed outcomes to differ from the function f that agents use in Periods 1 and 2 to evaluate their expected utility in Period 3. This is relevant in a number of cases. The function f^* may not take account of side-effects that make the treatment unattractive to the agent. Also, agents may not know the value of the average treatment effect that the researcher aims to estimate, and the assumptions they make about this in Periods 1 and 2

may be deviate systematically from the actual f^* .

The analysis in Section 3 describes how the choice of the effort level a and the validity of the ER depend on f. These decisions do not depend on f^* , so all results from that section also apply to the current framework.

To proceed, we start assuming that f satisfies (6), that effort is continuous, and that the actual outcomes Y are generated by $Y = f^*(X, a) + \varepsilon$ with

$$f^*(x,a) = \psi_0^* + \psi_1^* a + \psi_2^* x + \rho^* x a \tag{21}$$

The case $\psi_2 < \psi_2^*$ is particularly interesting as it captures situations where the agent perceives a disutility of the treatment that is not revealed in the outcome. Also, with $\rho < \rho^* < 0$, the agent dislikes the combination of high treatment and high effort more than is warranted by the corresponding actual outcome.

Equations (16) and (17), expressing the parameter of interest β and the probability limit of its estimator $\hat{\beta}_{IV}$ in terms of the model parameters, are now replaced by

$$\beta = \psi_2^* + \rho^* a_0 = \psi_2^* + \rho^* \frac{\psi_1 + \rho \overline{z}}{c_0}$$
(22)

$$\widehat{\beta}_{IV} \to \psi_2^* + \frac{\rho \psi_1^* + \rho^* \psi_1}{c_0} + \frac{\rho \rho^*}{c_0} \frac{\operatorname{cov}(Z, Z^2)}{\operatorname{var}(Z)}$$
(23)

In Appendix 5 we prove the following extension of Proposition 3,

Proposition 7. Consider the model with continuous effort, quadratic costs of effort, the agents' outcome utility function (6) with the conditions that ensure that it increases in effort and the treatment status, and the outcome function (21) with the conditions that (i) if $\rho^* \geq 0$ then $\psi_1^* \geq 0$, and (ii) if $\rho^* < 0$ then $\psi_1^* > 0$ and $Pr(Z < -\psi_1^*/\rho^*) = 1$. If the exclusion restriction is violated then the asymptotic bias of the IV estimator $\hat{\beta}_{IV}$ has the same sign as ρ , with the single exception that if $\rho^* = \psi_1^* = 0$ then this bias is zero.

The conditions (i) and (ii) on the range of ψ_1^* and the support of Z serve to ensure that $f^*(x, a)$ is increasing in a. This implies that we assume that both f and f^* are increasing in a. This is less restrictive as may seem. First, any disutility of a can to some extent be accommodated for by the cost function c(a)in Period 2. Secondly, we do not assume that the actual average treatment effect has the same sign as the perceived expected treatment effect, since we do not make assumptions on (the sign of) ψ_2^*, ρ^* , and ψ_0^* .

Proposition 7 implies that of the two interaction parameters, it is the interaction in the agent's objective function that drives the asymptotic bias. The underlying reason is that the sign of this interaction determines whether the agent's effort is increasing or decreasing in z. If it is increasing then this boosts the estimated average treatment effect regardless of how the actual outcomes generated. Notice also that the asymptotic bias is completely independent of the actual and the perceived additive treatment effect parameters ψ_2^* and ψ_2 and of the strength of the candidate instrument. The other propositions of Subsections 4.2 and 4.3 can be generalized accordingly.

It is sometimes plausible that the agent can have access to more information on the determinants of the treatment assignment process than the analyst can observe. As this is somewhat related to the topic of this subsection, we briefly discuss its main implications for the asymptotic bias here, using a simple framework where the assignment process can be captured by the following "treatment equation"

$$X = Z_1 + Z_2 + \omega \tag{24}$$

with $\mathbb{E}(\omega|Z) = 0$ and also $\mathbb{E}(\varepsilon|Z) = 0$, where $Z := (Z_1, Z_2)$. The analyst observes Y, X, and Z_2 , but not Z_1 . The agent is able to acquire his values of Z_1 and of Z_2 . We adopt the usual functional forms for f(x, a) and c(a) and we assume that $Z_1 \perp \!\!\perp Z_2$. Moreover, we define Z_1 and Z_2 such that $\mathbb{E}(Z_1) = 0$, so $\mathbb{E}(X|Z_2 = z_2) = z_2$. Notice that the parameter of interest β is now defined using Z_2 , as follows: $\beta := \operatorname{cov}(f(X, a_0), Z_2)/\operatorname{cov}(X, Z_2)$.

We can distinguish between four cases, depending on which Z_i are acquired by the agent. Firstly, suppose the ER applies to both Z_1 and Z_2 . Then either (or both) can be used as instrumental variables. Secondly, suppose the ER only applies to Z_1 . Then Z_1 is a valid instrument but it is unobserved to the analyst, so it cannot be used for inference. In this case, Z_1 can be subsumed into ω , which leads to the main model framework of this paper. Thirdly, suppose the ER only applies to Z_2 . Then Z_2 can be used for IV estimation.

Fourthly, and most interestingly, suppose the ER is violated for both Z_1 and Z_2 . By analogy to Subsection 4.1, we can rewrite β as $\psi_2 + \rho a_0$, where $a_0 = (\psi_1 + \rho \overline{z}_2)/c_0$. The IV estimator $\hat{\beta}_{IV}$ converges to $\operatorname{cov}(f(X, a(Z)), Z_2)/\operatorname{cov}(X, Z_2)$, where $a(z) = (\psi_1 + \rho(z_1 + z_2))/c_0$. This can be shown to lead to exactly the same expression for $\hat{\beta}_{IV}$ as (17) in Subsection 4.2, with Z_2 now replacing Z. Consequently, the results of the previous subsections all apply here. In sum, the results of the paper are robust with respect to whether the agent has more information on the treatment assignment process than the analyst. Note that the maintained assumption that the private information is orthogonal to the shared information (i.e. $Z_1 \perp Z_2$) is crucial.

4.6 Changing the variance of the candidate instrument or the randomization probability

By definition, it is difficult to empirically test the predictions of the paper, because exclusion restrictions are untestable. As an alternative approach, one may look for exogenous variation in the model parameters that affect the value of information. For a certain range of parameter values the ER is satisfied, whereas for another range it is not. With sufficient variation of these parameters in the data, it can be verified whether the estimated treatment effect behaves as predicted as a function of the parameter. Note that this does not involve a comparison between different values of a candidate IV. Rather, the comparison is at a deeper level, namely between different evaluation settings.

In fact, of all the determinants of V in the continuous case of Subsection 3.2, only the variance of Z is more or less directly observable. We therefore compare settings with different distributions of the candidate IV, notably with different variances. To proceed, let effort be continuous, and let Z have a distribution with separate parameters capturing the mean \overline{z} and the variance σ_z^2 . The data should now include settings with different values of σ_z^2 .

From the results in Subsection 3.2 it follows that the ER is satisfied iff $\sigma_z^2 < 2c_0\gamma/\rho^2$. For these (smaller) values of σ_z^2 , the IV estimator $\hat{\beta}_{IV}$ estimates the policy effect parameter β . This parameter depends on \overline{z} but not on σ_z^2 or any other feature of the distribution of Z. For the values of σ_z^2 larger than $2c_0\gamma/\rho^2$, the ER is violated, and $\hat{\beta}_{IV}$ estimates a number different from β , where the sign of the difference is determined by ρ . In general, this number varies itself with σ_z^2 . Moreover, $\hat{\beta}_{IV}$ as a function of σ_z^2 is discontinuous at $2c_0\gamma/\rho^2$.

In sum, the model predicts that as σ_z^2 increases, there is a point at which $\hat{\beta}_{IV}$ makes a discontinuous jump. This can be verified empirically. To the left of the discontinuity point, $\hat{\beta}_{IV}$ is constant, but this depends crucially on the usual functional forms for f(x, a) and c(a). If one has a priori knowledge on the sign of the interaction of x and a in f(x, a) then one can verify empirically whether the jump has the same sign. In practice, the verification of these predictions may be complicated by heterogeneity of agents in terms of other parameters (notably ρ, c_0 and γ) within and between evaluation settings with different σ_z^2 .

Now let us specialize the above setting by assuming that $Z \equiv X$ is binary, with $p := \Pr(Z = 1)$. This captures experiments in which participants can only observe their treatment status at a cost. The parameter p is then the randomization probability or treatment assignment probability. This is the only parameter of the distribution of Z, so it is not possible to vary $\operatorname{var}(Z)$ while keeping \overline{z} constant. Specifically, $\operatorname{var}(Z) = p(1-p)$ while $\mathbb{E}(Z) = p$. We therefore compare settings with different values of p which is directly observable.

The variance of Z is largest if p = 1/2. It follows that violation of the ER is most likely at values of p around 1/2. Specifically, the ER is satisfied iff $p(1-p) \leq 2c_0\gamma/\rho^2$, implying that if $\rho^2 < 8c_0\gamma$ then the ER is satisfied for all parameter values, whereas if $\rho^2 \geq 8c_0\gamma$ then the ER is violated if p is in a symmetric non-empty interval around p = 1/2. The latter interval is strictly embedded in (0, 1).

We first examine situations in which the ER is satisfied. Then the IV estimator $\hat{\beta}_{IV}$ estimates the policy effect parameter β , which equals

$$\beta = \psi_2 + \rho \frac{\psi_1 + \rho p}{c_0}$$

This parameter depends on p, meaning that the average treatment effect differs across experiments with different randomization probabilities, even if the ER is satisfied, like in a medical double-blind experiment. It is interesting to examine this in some detail. If p is high then the agents know that it is likely that they have been assigned to the treatment group. With $\rho > 0$ ($\rho < 0$), the marginal expected return of effort is higher (lower) if the expected treatment rate is higher, so agents then have an incentive to provide more (less) effort a_0 . For either sign of ρ , this boosts the treatment effect. For example, if p is large in a model with $\rho < 0$ (this could be labelled a moral hazard model), the outcome for non-treated agents is very low because all agents do not provide much effort, whereas the outcome for treated agents is high because of the treatment. If all agents would have provided sub-optimal medium effort then the outcome for treated agents would still be high, but the outcome for non-treated agents would be higher than before, resulting in a lower treatment effect. In sum, β increases in p for all parameter values.

We can relate these predictions to empirical findings from randomized experiments of the same medication with different randomization probabilities, and to empirical studies of randomized experiments in which participants were deliberately deceived about the randomization probability. This requires of course that the experiments are stretched over a sufficiently long time span for agents to be able to adjust their behavior in a way that affects the outcome. Also, the ER needs to be satisfied, so the time span should not be so long that agents can easily detect their treatment status.

Malani (2006) provides convincing evidence that the estimated treatment effect in blinded medical trials of ulcer medications and cholesterol medications increases with the announced treatment effect, and he attributes this result to placebo effects. In our framework this can be explained by agents rationally adjusting their efforts in response to the probability of being assigned to the treatment group. The trials of ulcer medications (H₂ blockers and proton-pump inhibitors) and cholesterol medications (statins) examined in Malani (2006) often last at least several weeks, and it is not inconceivable that agents change their lifestyle during the experiment in response to the randomization probability. For example, the treatment may be more effective if the agent simultaneously improves his lifestyle, and the agent will only do so if he is likely to be treated. This assumes $\rho > 0$, but one may also think of explanations based on $\rho < 0$. In fact, Malani (2006) also shows that in general the average outcome among treated increases stronger with p than the average outcome among controls, which is consistent with $\rho > 0$.

Penick and Hinkle (1964) conduct two types of experimental trials of an appetite-depressing drug (phenmetrazine). In all cases, p = 1/2. However, in the first trial type the agents are deceived, by, in effect, informing them that p = 0, whereas in the second type they are not deceived. Each trial type consists of about 25 daily events spread out over several months; on each of these days the participant receives the drug or placebo shortly before a standardized evening meal is served. The outcome is the amount of food not consumed at that meal. The estimated treatment effect on food intake reduction is larger in the second trial type than in the first, so again the treatment effect is larger if the perceived p is larger. Again, the outcome among treated increases more with the change in p than the outcome among controls, and this is consistent with $\rho > 0$. This can be explained if the amount of food intake on days between the trial event days (and in the morning of the trial event days) positively affects the extent to which one is able to have a reduced food intake when treated in the evening. After all, one may want to increase the food intake early in the day and on days between the trial events if one knows that p = 1/2 rather than p = 0.

We now turn to situations in which the ER is violated (so we adopt $\rho^2 \geq 8c_0\gamma$ and we consider values of p close to 1/2). The estimator $\hat{\beta}_{IV}$ estimates a number different from β , where the sign of the difference is determined by ρ . Specifically,

$$\widehat{\beta}_{IV} \to \psi_2 + \frac{2\rho\psi_1}{c_0} + \frac{\rho^2}{c_0}$$

It follows that as p increases from 0 to 1, there are two points at which $\widehat{\beta}_{IV}$ makes a discontinuous jump, and these jumps have opposite signs. This can be verified empirically (although, again, heterogeneity in other parameters may complicate this).¹⁵ If one has a priori knowledge on the sign of the interaction of x and a in f(x, a) then one can verify empirically that the sequence of the signs of the jumps is correct. For example, if the interaction sign is negative, so treatment and effort are substitutes, then one would expect the estimated treatment effect to be relatively small for treatment probabilities close to 1/2.

Instead of comparing evaluation settings with different observable distributions of Z, one might want to search for deeper instrumental variables affecting a specific unobserved model parameter, in order to assess predictions of the model framework concerning parameter regions in which the ER is satisfied or violated. However, if a parameter affects the estimated treatment effect in absence of the ER, then typically it also affects the treatment effect if the ER is satisfied, so an observed association between such instrumental variables and the estimated treatment effect is not informative on the predictions of the model framework (unless such an instrumental variable is so strongly related to the parameter that one can equate changes in one to changes in the other).

5 Conclusions

This paper takes the novel approach of using an economic framework to analyze the validity of the exclusion restriction and the properties of the IV estimator. The analysis has led to a large set of insights and implications for policy and treatment evaluation.

The exclusion restriction is more likely to be violated if the outcome of interest strongly depends on interactions between the agent's effort before the outcome is realized on the one hand, and the agent's treatment status on the other. We also find that the exclusion restriction is more likely to be violated if the candidate instrument covers a large shift in policy intensity or if it divides the population into groups of similar size.

Deliberate randomization of the intention to treat does not necessarily help. In social experiments, the randomization outcome is typically available to the agent at low cost. Moreover, the typical randomization probability of 1/2 corresponds to a high incentive to acquire one's realization. In fact, it may be better to use a smaller or larger probability, because this reduces the incentive to acquire and use the value of the candidate instrumental variable.

¹⁵Other predictions are that $\hat{\beta}_{IV}$ does not depend on p around p = 1/2 and that the discontinuity jumps are at values at equal distance from 1/2, but these results depends crucially on the functional forms for f(x, a) and c(a).

Having a weak instrument does not help either. Weakness of the candidate instrument, as defined or measured in ways proposed in the literature, is not directly related to the likelihood that the exclusion restriction is violated.

With discrete effort, instrumental variable analyses that restrict attention to agents with low or high resources (e.g. income) are more likely to be valid than analyses that include agents with intermediate resource levels. The reason is that for the former groups, the information is useless because it does not affect optimal effort. In addition, for the low resource agents, it may be too expensive.

In case of violation of the exclusion restriction, the sign and the magnitude of the bias are driven by the interaction term between the treatment status and the agent's effort in the outcome. The magnitude of the bias also depends on higherorder moments of the distribution of the candidate instrument. The strength of the candidate instrument and the size of the actual causal treatment effect do not affect the bias. In case of side effects and uncertainty about the actual effect, it is the interaction effect in the agent's objective function that drives the bias, rather than the interaction effect in the outcomes.

If the circumstances are such that it is plausible that the exclusion restriction is violated for the candidate instrumental variable, then IV cannot be applied. One can still use IV estimation to test the null hypothesis of a zero causal treatment effect, since under this hypothesis the exclusion restriction is not violated. However, to obtain effect estimates, one needs to look for other sources of exogenous variation in the treatment status, like random variation in the timing of the treatment, and impose a semi-parametric model structure (the so-called "timing-of-events" approach). In the paper we demonstrate that one may also exploit variation in the distribution of the candidate instrument across different evaluation settings. Alternatively, one may estimate a structural economic model.

The results also shed light on a number of other issues in the evaluation literature. For example, they provide a behavioral explanation of the positive relation between the randomization probability of a beneficial treatment and the associated treatment effect estimate. Such a positive relation has been observed in contexts in which a behavioral response is plausible. Further, as noted in the introduction, medical researchers have found associations between the methodological quality of empirical treatment evaluation studies and the treatment effect estimates found in those studies. Our findings concerning the determinants of the asymptotic effect bias can be used to understand such associations. If studies with inadequate concealment find larger effect estimates for a specific treatment, then this indicates that the treatment and the other effort towards the outcome are complements as outcome determinants.

Appendix

Appendix 1 Selective non-compliance

Suppose that agents can manipulate the probability distribution of their treatment status X, by way of choosing an action a before X is realized. As an extreme example, if the treatment status X is binary, they may switch treatment status by choosing an appropriate a. Agents' optimal a may depend on Z, which may be acquired at a cost.

This model framework can be reformulated in terms of our framework, by allowing a to affect the distribution of X|Z, which we denote by G(X|Z;a). If the agent does not acquire Z then he chooses action or effort a_0 leading to $G(X|Z;a_0)$ which we denote by H(X|Z). We start off by assuming that a does not affect the expected outcome function f, so we write f := f(X) instead of f(X,a). Note that this implies that the exclusion restriction is satisfied.

To see the connection to our framework, note that determining the optimal a involves calculation of $\mathbb{E}_{X|Z=z;a}f(X)$. This equals $\int f(x)dG(x|z;a)$, which can be rewritten as¹⁶

$$\int f(x)dG(x|z;a) = \int \left[\frac{f(x)dG(x|z;a)}{dH(x|z)}\right] dH(x|z)$$

The term in square brackets can now be interpreted as a new outcome function \tilde{f} and we can use the results derived earlier in the section to study the extent to which agents acquire their value of z.

Under certain conditions, IV estimation accommodates selective non-compliance, where it is (often tacitly) assumed that the compliance decision is made upon the realization of the treatment status X. In the setup of this subsection, the compliance decision is made *before* the realization of X, but in this setup the ER is not violated either, because the actual outcome f only depends on X. Obviously, the setup of this subsection and the framework earlier in the paper can be combined by allowing the original f to depend on a as well.

Appendix 2 Proof of Proposition 3 and some implications

First of all, recall that $Pr(Z \ge 0) = 1$ and var(Z) > 0. From (18), if $\rho = 0$ then the asymptotic bias is zero.

¹⁶If z is a location parameter of the distribution of [X|Z = z; s], so that G(X|z; a) can be expressed as $G_0(X - z|a)$ with G_0 functionally independent of z, then this can be further simplified to $\int f(x+z) dG_0(x|a)$.

Now consider the case $\rho > 0$. From Subsection 3.2, the conditions on the outcome function that are listed in the proposition amount to $\psi_0 > 0, \psi_1 \ge 0, \psi_2 \ge 0$. Also, $c_0 > 0$. By substituting $\operatorname{cov}(Z, Z^2) = \mathbb{E}(Z^3) - \overline{z}\mathbb{E}(Z^2)$ into (18), it follows that asymptotically $\widehat{\beta}_{IV} - \beta$ has the same sign as

$$\psi_1 + \rho \overline{z} \left[\frac{\widetilde{\mu}_3 - 2\widetilde{\mu}_2 + 1}{\widetilde{\mu}_2 - 1} \right] \tag{25}$$

where $\tilde{\mu}_i := \mathbb{E}Z^i/\overline{z}^i$ for i = 2, 3. Note that the denominator of the above term in square brackets is proportional to $\operatorname{var}(Z)$ which is positive. Consequently, the term in square brackets is positive iff the numerator is positive. The latter is equivalent to $\tilde{\mu}_3 > \tilde{\mu}_2^2 - (\tilde{\mu}_2 - 1)^2$. From Shohat and Tamarkin (1943)'s results for the so-called Stieltjes Moment Problem it follows that a random variable Zwith $\operatorname{Pr}(Z \ge 0) = 1$ and $\operatorname{var}(Z) > 0$ necessarily satisfies $\mathbb{E}(Z)\mathbb{E}(Z^3) > (\mathbb{E}(Z^2))^2$, with the exception of the special case in which the support of Z consists of two mass points, one of which is zero. The inequality is equivalent to $\tilde{\mu}_3 > \tilde{\mu}_2^2$, which in turn implies that $\tilde{\mu}_3 > \tilde{\mu}_2^2 - (\tilde{\mu}_2 - 1)^2$. For the special case in which the support of Z consists of two mass points, one of which is zero, there holds that $\tilde{\mu}_3 = \tilde{\mu}_2^2$, but also that $\tilde{\mu}_2 - 1 > 0$, so that again $\tilde{\mu}_3 > \tilde{\mu}_2^2 - (\tilde{\mu}_2 - 1)^2$. Consequently, the asymptotic bias is always positive.

Now consider the case $\rho < 0$. From Subsection 3.2, the conditions on the outcome function that are listed in the proposition amount to $\psi_0 > 0, \psi_1 > 0, \psi_2 > 0, \Pr(X < \frac{\psi_1}{-\rho} | Z = z) = 1$ for all possible realizations z, and $\psi_2 c_0 + \rho \psi_1 > 0$, with $c_0 > 0$. From equation (18), and by analogy to equation (25), we obtain that asymptotically $\hat{\beta}_{IV} - \beta$ has the same sign as

$$\frac{\psi_1}{\rho} + \frac{\mu_3 - 2\overline{z}\mu_2 + \overline{z}^3}{\mu_2 - \overline{z}^2} \tag{26}$$

where $\mu_i := \mathbb{E}(Z^i)$. In Subsection 3.2 we saw that the parameter inequalities for this case imply that $\Pr(Z < \frac{\psi_1}{-\rho}) = 1$, so Z satisfies $\Pr(Z \in [0, -\psi_1/\rho)) = 1$. We denote $-\psi_1/\rho$ by z_u . The term (26) is negative iff

$$-z_u(\mu_2 - \overline{z}^2) + \mu_3 - 2\overline{z}\mu_2 + \overline{z}^3 \tag{27}$$

is negative. For convenience, we rewrite this in terms of central moments of Z. Let $\kappa_i := \mathbb{E}(Z - \overline{z})^i$. Then (27) equals

$$\kappa_3 - (z_u - \overline{z})\kappa_2 \tag{28}$$

To sign this, we apply results for the so-called Hausdorff Moment Problem (see Shohat and Tamarkin, 1943, and Frontini and Tagliani, 1995). In particular, a random variable Z_0 with $\Pr(Z_0 \in [0, 1)) = 1$ and $\operatorname{var}(Z_0) > 0$ necessarily satisfies

$$(1 - \mathbb{E}(Z_0))(\mathbb{E}(Z_0^2) - \mathbb{E}(Z_0^3)) > (\mathbb{E}(Z_0) - \mathbb{E}(Z_0^2))$$

We take $Z = z_u \cdot Z_0$, so that Z satisfies $\Pr(Z \in [0, z_u)) = 1$ and $\operatorname{var}(Z) > 0$, as required, and we use the notation \overline{z} and μ_i to denote its moments. Clearly, $\overline{z} = z_u \mathbb{E}(Z_0)$ and $\mu_i = z_u^i \mathbb{E}(Z_0^i)$. The above moment inequality can now be written as

$$z_u^2(\mu_2 - \overline{z}) - z_u(\mu_3 - \overline{z}\mu_2) + \overline{z}\mu_3 - \mu_2^2 > 0$$

and in terms of central moments of Z this simplifies to

$$0 > \kappa_2^2 + (z_u - \overline{z}) \left[\kappa_3 - (z_u - \overline{z}) \kappa_2 \right]$$

Since $z_u > \overline{z}$, this implies that $\kappa_3 - (z_u - \overline{z})\kappa_2 < 0$. Consequently, the asymptotic bias is always negative. This completes the proof of Proposition 3.

Notice that the result for $\rho < 0$ also applies if we allow for X and Z to be able to attain the value z_u . Also, notice that by substituting $\operatorname{cov}(Z, Z^2) = \mathbb{E}(Z^3) - \overline{z}\mathbb{E}(Z^2)$ into (18) and by rewriting this in terms of central moments κ_i of Z, we obtain

$$\widehat{\beta}_{IV} - \beta \to \frac{\rho \psi_1}{c_0} + \frac{\rho^2 \overline{z}}{c_0} \left[\frac{\kappa_3}{\overline{z} \kappa_2} + 1 \right]$$
(29)

By substituting β into the right-hand side, and by noting that $\kappa_3/(\bar{z}\kappa_2)$ equals the skewness of Z times the coefficient of variation of Z, we obtain equation (19) in the main text. Also, by combining equation (29) with the proof for the case $\rho > 0$, it becomes clear that in that case $\hat{\beta}_{IV}$ and $\hat{\beta}_{IV} - \beta$ asymptotically increase in ρ .

Appendix 3 Proof of Proposition 4

Because c(a) is not parameterized, there are no parameterized expressions for a(z) and a_0 either. However, from equations (3) and (4) it follows that with the outcome function (6) and with the assumptions on f and on c(a) listed in the proposition, the optimal effort level a(z) in absence of the ER increases in z (whereas with $\rho < 0$ it decreases in z). Moreover,

$$c'(a_0) = \psi_1 + \rho \overline{z} = c'(a(\overline{z}))$$

implying that $a(\overline{z}) = a_0$. The latter means that the agent's optimal effort if he knows that he has the average treatment intensity is equal to the agent's optimal effort if he does not know his intensity. We will use this below. Note that the function a(z) cannot be constant over the support of Z.

If we substitute (6) into (13) then we obtain that

$$\beta = \psi_2 + \rho a_0 \tag{30}$$

Similarly, if we substitute (6) into (14) then we obtain that

$$\widehat{\beta}_{IV} \to \psi_2 + \frac{\psi_1 \text{cov}(Z, a(Z))}{\text{var}(Z)} + \frac{\rho \text{cov}(Z, Za(Z))}{\text{var}(Z)}$$
(31)

Since a(z) is increasing, the second term on the right-hand side is positive. By comparing (31) to (30), it then follows that if the third term on the right-hand side of (31) is larger than or equal to ρa_0 then asymptotically $\hat{\beta}_{IV} > \beta$. Since $\rho > 0$, this requires that

$$\frac{\operatorname{cov}(Z, Za(Z))}{\operatorname{var}(Z)} \ge a_0$$

which is equivalent to $cov(Z, Z[a(Z) - a_0]) \ge 0$. The latter can be written as

$$\mathbb{E}(Z^2 a(Z)) - \overline{z} \mathbb{E}(Z a(Z)) - a_0 \mathbb{E}(Z^2) + a_0 \overline{z}^2 \ge 0$$

which is equivalent to

$$\mathbb{E}\left[Z(Z-\overline{z})(a(Z)-a_0)\right] \ge 0 \tag{32}$$

Earlier in this proof we derived that a(z) is increasing and that $a(\overline{z}) = a_0$. These two facts imply that $a(z) \ge a_0 \iff z > \overline{z}$. Consequently, (32) is true, and this completes the proof. \Box

Note that we can rewrite equation (31) as

$$\widehat{\beta}_{IV} \to \beta + \frac{\psi_1 \text{cov}(Z, a(Z))}{\text{var}(Z)} + \frac{\rho \mathbb{E}\left[Z(Z - \overline{z})(a(Z) - a_0)\right]}{\text{var}(Z)}$$
(33)

Appendix 4 Proof of Proposition 5

In terms of our model and notation, Athey and Levin (2001) prove that the assumptions on the function f imply that the optimal effort level a(z) in absence of the ER increases in z.

The numerator on the right-hand side of (15), $\operatorname{cov}(\mathbb{E}(f(X, a(Z)) - f(X, a_0)|Z), Z)$ determines the sign of the asymptotic bias. Following the line of reasoning that leads up to (32) in Appendix 3, it is easy to see that this numerator can be written as

$$\mathbb{E}\left[(Z - \overline{z}) \cdot \mathbb{E}(f(X, a(Z)) - f(X, a_0)|Z)\right]$$
(34)

Since a increases, $a(z) \ge a(\overline{z}) \iff z \ge \overline{z}$. Now f is non-negative, increasing and supermodular, so for any given x there holds that $f(x, a(z)) \ge f(x, a(\overline{z}))$ if and only if $z \ge \overline{z}$. By assumption, $a(\overline{z}) = a_0$, so for any given x there holds that $f(x, a(z)) \ge f(x, a_0)$ if and only if $z \ge \overline{z}$. The expectation over X|Z = zof $f(X, a(z)) - f(X, a_0)$ then also has the property that it is non-negative if and only if $z \ge \overline{z}$. Consequently, (34) is non-negative, and this completes the proof.

Appendix 5 Proof of Proposition 7

Equations (22) and (23) imply that

$$\widehat{\beta}_{IV} - \beta \to \frac{\rho}{c_0} \left[\psi_1^* - \rho^* \overline{z} + \rho^* \frac{\operatorname{cov}(Z, Z^2)}{\operatorname{var}(Z)} \right]$$
(35)

Clearly, if $\rho = 0$ or if $\rho^* = \psi_1^* = 0$ then the asymptotic bias is zero.

Now consider the other cases, so $\rho \neq 0$, and, in addition, $\rho^* \neq 0$ and/or $\psi_1^* > 0$. The proof in Appendix 2 of Proposition 3 demonstrates that if $\rho^* > 0$ then

$$\psi_1^* - \rho^* \overline{z} + \rho^* \frac{\operatorname{cov}(Z, Z^2)}{\operatorname{var}(Z)}$$
(36)

is positive. If $\rho^* = 0$ then $\psi_1^* > 0$ and again (36) is positive. Similarly, the proof in Appendix 2 demonstrates that if $\rho^* < 0$ then the expression in (36) divided by ρ^* is negative, taking into account the assumptions that then $\psi_1^* > 0$ and $\Pr(Z < -\psi_1^*/\rho^*) = 1$. So if $\rho^* < 0$ then again (36) is positive. By comparing (35) to (36) it follows that the asymptotic bias has the same sign as ρ . Note that we do not require assumptions on the range of values of ψ_0^* and ψ_2^* and the support of X|Z in terms of $-\psi_1^*/\rho^*$ that correspond to those made concerning the parameters of f. \Box

References

- Abbring, J.H. and G.J. van den Berg (2003), "The non-parametric identification of treatment effects in duration models", *Econometrica*, 71, 1491–1517.
- Abbring, J.H. and G.J. van den Berg (2005), "Social experiments and instrumental variables with duration outcomes", Working paper, Free University Amsterdam, Amsterdam.
- Angrist, J.D., G.W. Imbens, and D.B. Rubin (1996), "Identification of causal effects using instrumental variables", *Journal of the American Statistical As*sociation, 91, 444–455.
- Athey, S. and J. Levin (2001), "The value of information in monotone decision problems", Working paper, Stanford University, Stanford.
- Blundell, R. and T. MaCurdy (1999), "Labor supply", in O. Ashenfelter and D. Card, editors, *Handbook of Labor Economics, Volume III*, North-Holland, Amsterdam.
- Blundell, R. and J.L. Powell (2003), "Endogeneity in nonparametric and semiparametric regression models", in M. Dewatripont et al., editors, Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress, Vol. II, Cambridge University Press, Cambridge.
- Chan, T.Y. and B.H. Hamilton (2006), "Learning, private information and the economic evaluation of randomized experiments", *Journal of Political Econ*omy, 114, 997–1040.
- Epstein, S. (1996), Impure Science: AIDS, Activism, and the Politics of Knowledge, University of California Press, Berkeley.
- Frontini, M. and A. Tagliani (1995), "Maximum entropy in the finite Hausdorff moment problem", Working paper, Politecnico di Milano, Milan.
- Gollier, C. (2001), The Economics of Risk and Time, MIT Press, Cambridge.
- Hansen, L.P. and T.J. Sargent (1980), "Dynamic linear rational expectations models", Journal of Economic Dynamics and Control, 2, 7–46.
- Heckman, J.J., R.J. LaLonde, and J.A. Smith (1999), "The economics and econometrics of active labor market programs", in O. Ashenfelter and D. Card, editors, *Handbook of Labor Economics, Volume III*, North-Holland, Amsterdam.
- Imbens, G.W. and J.D. Angrist (1994), "Identification and estimation of local average treatment effects", *Econometrica*, 62, 467–475.

- Malani, A. (2006), "Identifying placebo effects with data from clinical trials", Journal of Political Economy, 114, 236–256.
- Penick, S.B. and L.E. Hinkle (1964), "The effect of expectation on response to phenmetrazine", *Psychosomatic Medicine*, 26, 369–373.
- Persico, N. (2000), "Information acquisition in auctions", *Econometrica*, 68, 135–148.
- Philipson, T. and J. DeSimone (1997), "Experiments and subject sampling", *Biometrika*, 84, 619–631.
- Robins, J.M. and S. Greenland (1996), "Identification of causal effects using instrumental variables: comment", *Journal of the American Statistical Association*, 91, 456–458.
- Rosenzweig, M.R. and K.I. Wolpin (2000), "Natural "natural experiments" in economics", *Journal of Economic Literature*, 38, 827–874.
- Schuklenk, U. (2003), "AIDS: bioethics and public policy", New Review of Bioethics, 1, 127–144.
- Schulz, K.F. (1995), "Subverting randomization bias in controlled trials", Journal of the American Medical Association, 274, 1456–1458.
- Schulz, K.F., I. Chalmers, R.J. Hayes, and D.G. Altman (1995), "Empirical evidence of bias: dimensions of methodological quality associated with estimates of treatment effects in controlled trials", *Journal of the American Medical* Association, 273, 408–412.
- Schulz, K.F. and D.A. Grimes (2002), "Allocation concealment in randomised trials: defending against deciphering", *Lancet*, 359, 614–618.
- Shohat, J.A. and J.D. Tamarkin (1943), *The Problem of Moments*, American Mathematical Society, New York.
- Stock, J.H., J.H. Wright, and M. Yogo (2002), "A survey of weak instruments and weak identification in generalized method of moments", *Journal of Business* and Economic Statistics, 20, 518–529.