Learning in Games and the Interpretation of Natural Experiments[1]

Drew Fudenberg[2] and David K. Levine[3]

Abstract

We examine natural experiments where the variable of interest is the effort levels of the agents, the treatment and control correspond to success or failure, and there is unobserved heterogeneity in the agents' efforts. We show that in such experiments the treatment effect estimated by standard methods such as regression discontinuity analysis or difference-in-difference may contain a transient "learning effect" that is entangled with the structural long-term effect of the treatment. This learning effect occurs when agents are uncertain of the effectiveness of their effort: their success or failure signals agents how much their effort matters to success, and consequently changes the amount of effort they provide after treatment. We examine which direction the effect goes and when it is likely to be present. We illustrate this with applications taken from the literature, and show how under some circumstances the presence of learning can alter policy conclusions.

[1]We thank Chris Ackerman, Josh Angrist, Daniel Clark, Michael Griffiths, Andrea Ichino, Andrea Mattozzi, and Parag Pathak for helpful comments and conversations. We are grateful to NSF grant SES 1643517 and the EUI Research Council for financial support.

[2]Department of Economics, MIT; drew.fudenberg@gmail.com

[3]Department of Economics and RSCAS, EUI, and Department of Economics, WUSTL; david@dklevine.com

1. Introduction

Natural experiments have become widely used in economics as a way of estimating treatment effects in non-experimental settings. This paper points out an important issue to consider in the construction and interpretation of these experiments: When selection into the treatment depends upon the subjects' effort, the treatment effect may depend on their state of knowledge, and this may vary over time and place. Thus the estimated treatment effect may not be as robust or causal as might have been supposed.

We focus on natural experiments where the "treatment" corresponds to "success" or "failure." Being admitted to an exclusive school is a success, being bombed by the government is a failure, winning an election is a success. In many of these circumstances success or failure depends upon effort. While agents undoubtedly know this, they are generally not certain of the extent to which their effort matters. In such settings, if the agent makes an effort it is natural to interpret success as indicating "effort does indeed matter" and failure as indicating "effort does not matter so much." Our starting point is a simple model of Bayesian updating in which success indeed signals effort is more likely to matter, and failure signals it is less likely to matter.

Subjects' uncertainty about how much effort matters has an important consequence for the evaluation of treatment effects: If a success signals that effort matters, then success will lead to increased effort, and if failure signals effort does not matter, then failure will lead to decreased effort. Hence the consequence of the treatment has two parts: a direct effect on preferences and/or technology and a learning effect. Moreover, the learning effect has a particular direction: the treatment corresponding to "success" leads to better outcomes due to increased effort, and so the estimated treatment effect will be higher than it would be in the absence of learning.

From the point of view of policy the relative importance of the direct effect and the learning effect makes a difference: how much learning takes place depends upon what is known in advance, so the effect of the treatment depends upon the state of knowledge. This in turn depends upon the amount of information available in advance, upon the amount of experience the agents have and so forth. To take a particular example: if being admitted to an exclusive school results in better outcomes one might conclude that random admissions would be a fair way of distributing these benefits. Suppose, however, the better outcomes are due to learning and greater effort. Once it is known that admissions are random, outcomes will no longer signal how much effort matters. Consequently admission to an exclusive school will no longer lead to increased effort and better outcomes: the treatment effect will disappear and the policy will be a failure.

We develop precise characterizations of the bias that can arise from neglected learning effects in two leading methods for analyzing natural experiments, namely regression discontinuities and difference in differences.[4] We then apply our analysis to two well-known natural experiments, those of Dell and Querubin [2] and Lee [4]. Dell and Querubin [2] uses a novel data set based on the indices used by the US Air Force to decide which hamlets to bomb during the Vietnam War. The paper carefully and convincingly argues that the hamlets that were bombed worked less hard to accommodate American interests. That finding makes it tempting to conclude that the bombing campaign was counterproductive, but we show that this does not follow. Intuitively, the idea is that the threat of bombing may induce compliance, while being bombed may convince villagers that compliance is pointless as they will be bombed anyway. Lyall [5] studies the indiscriminate shelling of Chechen villages by the Russians, where this learning effect is less likely to be present, which may be why the paper reaches the opposite conclusion from Dell and Querubin [2]. This raises the possibility that the preference effect was positive, but the learning effect was negative. Because the econometrician only observes the aggregate effect, the coefficient estimate must be interpreted with care..

Lee [4] studies the effect of incumbency in elections to the U.S. House of Representatives, and finds that winning an election improves future election prospects relative to losing. The paper interprets this as showing that incumbents have an electoral advantage, perhaps due to a lower cost of campaigning. In principle this finding is also subject to our criticism, as it might be that winning signals that campaign effort is effective and so leads to greater effort in future elections. However in this context we argue that the learning effect is likely to be small. In particular the candidates not only observe success or failure, but also the vote differential and under these circumstances success and failure provide little additional information about the benefits of effort.

Section 2 of the paper lays out our model of agent behavior. Section 3 analyses the model and establishes the comparative statics that we use in our discussion of how to interpret the econometric findings. Section 4 computes the implications of our model for the estimated treatment effects in regression discontinuity and difference-in-difference econometric specifications. It shows that the estimated treatment effects are the sum of a preference effect that is independent of the subjects' information and a learning effect that depends on that information. Section 5 then applies our findings to the study of state violence and incumbency.

[4] See Angrist and Pischke [1] for example.

2. The Agent's Problem

We study a simple two period learning problem of a single agent. In each period $t = 1, 2$, the agent determines a level of effort $e_t \in [0, 1]$. At the end of each period, a random variable z_t with continuous density determines a payoff relevant outcome $X_t \in \{0, 1\}$, where $X_t = 0$ if $z_t \leq 0$ represents "failure" and $X_t = 1$ if $z_t > 0$ represents "success"". In the empirical setting X_t will constitute the *treatment* whose effectiveness the econometrician is trying to ascertain. Regardless of the outcome X_t, effort e_t provides a benefit $u(e_t)$, where u is smooth with $u' > 0, u'' < 0$ and $\lim_{x \to 0} u'(x) = \infty$. In addition, the agent obtains an additive bonus of 1 if $X_t = 1$. Exerting effort in period 1 has constant marginal cost of $c_1 \geq \underline{c}$, in the first period, and constant marginal cost of $c_2(X_1)$ in the second period which may depend on the first period's outcome.[5] In addition to its direct benefit, effort can influence the probability of a success. To model this, we assume that the cumulative distribution of the unfavorable outcome in period t is $F(-\gamma e_t)$, with $1 \geq \gamma \geq 0$, so that γ parameterizes the effectiveness of effort.

One way the first period outcome influence second period effort is through its effect on the agent's second period cost of effort. Our focus is on a second channel, namely the effect of the first period outcome on the agent's beliefs about γ. To model the agent's inference process, we assume that success or failure is determined by underlying score or index variables z_1 and z_2 which are independent, with cumulative distribution function F, so that exerting maximal effort of 1 causes an additive shift of the distribution of z by γ, and exerting a fraction of the maximal effort shifts the distribution proportionally. Moreover we assume that in the range $[-1, 0]$the cdf $F(z_t)$ is smooth with $F''(z_t) > 0$ and $F'(z) > -zF''(z)$.[6] The agent knows the function F, and that the probability of success is $1 - F(-\gamma e)$, but is uncertain about the value of γ, which measure the effectiveness of effort.

For simplicity we assume that the agent contemplates two possibilities: $1 \geq \gamma = \overline{\gamma} > \underline{\gamma} \geq 0$, and assigns prior probability $1 > p_1 > 0$ to the state of the world $\gamma = \overline{\gamma}$ where effort matters more. Finally, as a technical assumption, we assume that cost is not too low relative to the marginal benefit of effort: $\underline{c} > \overline{\gamma} F'(-\overline{\gamma}) + u'(1)$.

We assume that the agent observes the outcome X_1 at the end of the first period, but no other information about the effectiveness of effort. Let $p_2(X_1)$ denote the agent's posterior after observing X_1, as determined by Bayes law. Note that when

[5]Nothing of importance would change if the second period cost were a stochastic function of the first period outcome.

[6]This is satisfied, for example, if the cdf has the form $F_0 \exp(z/\sigma)$ for $\sigma \geq 1$.

$e = 0$ success or failure conveys no information about γ. More generally as we show in Theorem 2 below higher levels of e generate more information about γ than lower ones, at least when $\underline{\gamma}$ is small.

We assume that the agent is myopic and chooses e_t solely to maximize utility in period t. This means that in period 2 they use the information from the outcome in period 1, but do not invest in setting e_1 in order to have a more informative signal about how much effort matters. Implicitly, this also means that the time elapsed between effort and the resulting score and outcome is much less than that between the periods.

Hence the objective function for the player in a given period is

$$v(e_t|p_t, c_t) = p_t(1 - F(-\overline{\gamma}e_t)) + (1 - p_t)(1 - F(-\underline{\gamma}e_t)) + u(e_t) - c_t e_t.$$

In this formulation, $c_2 - c_1$captures all of the changes in preferences and technology brought about by the treatment X_1,so it implicitly assumes that all such changes are proportional to e.

We will consider two applications. One is a setting of civil strife. where a central government faces an insurgency and bombs or shells agents who correspond to villages. Here effort corresponds to resisting the insurgents and the treatment corresponds to being bombed which is $X_1 = 0$. . One theory of the effect this might have is that bombing angers villagers and makes them less sympathetic to the government and more sympathetic to the insurgents, so that villagers are less inclined to resist. This corresponds to a change in preferences in which $c_2(0) > c_2(1)$: bombing lowers the marginal utility of effort. The other theory is that bombing destroys infrastructure and makes the village less attractive to insurgents. In this case it is easier for villagers to resist. This corresponds to a reduction in the cost of resistance so that $c_2(0) < c_2(1)$.

In the other application the setting is an election, the agent is a candidate, and the outcome $X_1 = 1$ is winning the election while $X_1 = 0$ is losing the election. Here the treatment corresponds to winning $X_1 = 1$. In this case the theory is that winning the election gives the winner, the incumbent, an advantage in subsequent election which is the same as lowering the cost of effort. That is, $c_2(1) < c_2(0)$. Alternatively we could hypothesize that incumbency raises the cost of effort. This would be the case if personal campaigning is important and the politician who lives and works in Washington D.C. finds it costly to campaign in her local district.

Note that the model assumes that the agent knows the threshold for a good outcome, and is only uncertain of how much effort influences the probability that this occurs, with the probability of success for a given effort level independent across periods. If the agent were also uncertain about the cutoff levels, they would face a more complicated inference problem.

3. Solution of the Agent's Problem

For given values of c_t and p_t the agent faces a simple static optimization problem with objective function $v(e_t|p_t, c_t)$. Under our assumptions this has a unique solution $\hat{e}(p_t, c_t)$. Our first result gives the key properties of that solution:

Theorem 1. *The solution $\hat{e}$ of the agent's problem is the unique solution to*

$$p_t \left[\overline{\gamma} F'(-\overline{\gamma} e_t) - \underline{\gamma} F'(-\underline{\gamma} e_t)\right] + \underline{\gamma} F'(-\underline{\gamma} e_t) + u'(e_t) = c_t. \tag{1}$$

The solution is strictly interior, strictly increasing in p_t, and strictly decreasing in c_t. Moreover, there is a $\gamma^* \in (\underline{\gamma}, \overline{\gamma})$ such that

$$\frac{\partial \hat{e}_t}{\partial p_t} = \frac{[\gamma F'(\gamma \hat{e})]'(\gamma^* \hat{e})}{-v''(\hat{e}, p_t)} (\overline{\gamma} - \underline{\gamma}). \tag{2}$$

Proof. Since $1 \geq \gamma = \overline{\gamma} > \underline{\gamma} > 0$, the agent's effort choice can vary the probability of a success from $1 - F(0)$ to $1 - F(-1)$. Over this range, the second derivative of the agent's objective function is

$$-p_t \overline{\gamma}^2 F''(-\overline{\gamma} e_t) - (1 - p_t) \underline{\gamma}^2 F''(-\underline{\gamma} e_t) + u''(e_t)$$

in which all terms are negative by assumption. Thus the optimum is unique. The assumption that $\lim_{x \to 0} u'(x) = \infty$ forces $\hat{e} > 0$, while $\underline{c} > \overline{\gamma} F'(-\overline{\gamma}) + u'(1)$ forces $\hat{e} < 1$.

To prove 2 we apply the implicit function theorem to the first order condition (1). The derivative of the first order condition with respect to c_t is negative so $\partial \hat{e} / \partial c < 0$. The derivative with respect to p_t is $\overline{\gamma} F'(-\overline{\gamma} e_t) - \underline{\gamma} F'(-\underline{\gamma} e_t)$. This is strictly positive from our assumption that $F'(z) > -zF''(z)$, a fact that we will use also in proving Theorem 2. Hence $\partial \hat{e} / \partial p_t > 0$.

The final expression follows from the implicit function theorem and the mean value theorem. □

Our second main result is that the posterior probability that effort matters "more", that is, that $\gamma = \overline{\gamma}$, is higher after success, and that this effect is weaker when the agent's prior is closer to 0 or 1 as measured by $\kappa_1 = |1 - 2p_1|$. We also show that higher levels of e_1 are more informative when $\underline{\gamma}$ is small . We will use these results in the next section to analyze the sensitivity of the estimated treatment effects to the confounding effects of the agent's inference process.

Theorem 2. *$p_2(X_1)$ is strictly increasing in X_1 with*

$$|p_2(X_1) - p_1| \leq \frac{1}{\min\{F(-\overline{\gamma}), 1 - F(0)\}}(1 - \kappa_1),$$

$p_2(0)$ is strictly decreasing in e_1, and for $\underline{\gamma}$ sufficiently small $p_1(1)$ is strictly increasing in e_1.

Proof. For any $e_1 > 0$, the agent's posterior probability that $\gamma = \overline{\gamma}$ after observing $X_1 = 0$ is

$$\begin{aligned} p_2(0) &= \left(\frac{F(-\overline{\gamma}e_1)}{p_1 F(-\overline{\gamma}e_1) + (1 - p_1)F(-\underline{\gamma}e_1)} \right) p_1 \\ &= \left(\frac{1}{p_1 + (1 - p_1)F(-\underline{\gamma}e_1)/F(-\overline{\gamma}e_1)} \right) p_1. \end{aligned}$$

This will be strictly decreasing in e_1 provided $F(-\underline{\gamma}e_1)/F(-\overline{\gamma}e_1)$ is strictly increasing. To see that this is the case, differentiate $F(-\underline{\gamma}e_1)/F(-\overline{\gamma}e_1)$ to find that the sign of its derivative is determined by

$$\begin{aligned} &\overline{\gamma}F'(-\overline{\gamma}e_1)F(-\underline{\gamma}e_1) - \underline{\gamma}F'(-\underline{\gamma}e_1)F(-\overline{\gamma}e_1) \\ &> F(-\underline{\gamma}e_1)\left[\overline{\gamma}F'(-\overline{\gamma}e_1) - \underline{\gamma}F'(-\underline{\gamma}e_1)\right]. \end{aligned}$$

As we showed in the proof of Theorem 1 $\overline{\gamma}F'(-\overline{\gamma}e_1) - \underline{\gamma}F'(-\underline{\gamma}e_1) > 0$ which gives the desired result.

The posterior after observing $X_1 = 1$is

$$\begin{aligned} p_2(1) &= \frac{1 - F(-\overline{\gamma}e_1)}{p_1(1 - F(-\overline{\gamma}e_1)) + (1 - p_1)(1 - F(-\underline{\gamma}e_1))} p_1 \\ &= \frac{1}{p_1 + (1 - p_1)(1 - F(-\underline{\gamma}e_1))/(1 - F(-\overline{\gamma}e_1))} p_1. \end{aligned}$$

For $\underline{\gamma} = 0$ the derivative with respect to e_1 is strictly positive so this remains the case for all sufficiently small $\underline{\gamma}$.

Since F is strictly increasing, $F(-\underline{\gamma}e_1)/F(-\overline{\gamma}e_1) > 1$ and $(1 - F(-\underline{\gamma}e_1))/(1 - F(\overline{\gamma}e_1)) < 1$. Thus $p_2(0) < p_1$and $p_2(1) > p_1$, so $p_2(X_1)$ is strictly increasing in X_1.

To prove the second claim in the theorem, let $\Pr(X_1|\gamma)$ denote the conditional probability of the outcome X_1 given γ and the agent's optimal choice of first-period effort$\hat{e}_1$, so that $\Pr(1|\gamma) = 1 - F(-\gamma\hat{e}_1)$ and $\Pr(0|\gamma) = F(-\gamma\hat{e}_1)$. Then Bayes Law gives

$$p_2(X_1) = \left(\frac{\Pr(X_1|\overline{\gamma})}{p_1 \Pr(X_1|\overline{\gamma}) + (1-p_1)\Pr(X_1|\underline{\gamma})} \right) p_1$$

so

$$\left| \frac{p_2(X_1) - p_1}{p_1} \right| = \left| \frac{(1-p_1)(\Pr(X_1|\overline{\gamma}) - \Pr(X_1|\underline{\gamma}))}{p_1 \Pr(X_1|\overline{\gamma}) + (1-p_1)\Pr(X_1|\underline{\gamma})} \right|$$
$$\leq \left| \frac{1}{\min\{\Pr(X_1|\overline{\gamma}), \Pr(X_1|\underline{\gamma})\}} \right|.$$

The result then follows from minimizing the probabilities in the denominator over $\hat{e}_1$.

By reversing the role of $\underline{\gamma}$ and $\overline{\gamma}$ the bound also holds for $1-p_1$. Since $\min\{p_1, 1-p_1\} = 1 - |1-2p_1|$ the result follows. □

4. The Econometrician's Problem

Now we turn to the problem faced by the econometrician. The econometrician observes a sample of different agents indexed by "location" i, with parameters c_{1i} and p_{1i}, and would like to determine the effect of X_{1i} on $\hat{e}_{2i}$. This is called the treatment effect, and is defined as

$$D_i = \hat{e}_{2i}(p_{2i}(1), c_{2i}(1)) - \hat{e}_{2i}(p_{2i}(0), c_{2i}(0)).$$

Under modest assumptions this treatment effect is independent of i. First, we assume that $\hat{e}(p_t, c_t)$ is approximately linear,[7] that is$b = \partial\hat{e}_t/\partial c_t$ and $\partial\hat{e}_t/\partial p_t$ are approximately constant, where from Theorem 1

$$\frac{\partial \hat{e}_t}{\partial p_t} = \frac{[\gamma F'(\gamma \hat{e}_t)]'(\gamma^* \hat{e}_t)}{-v"(\hat{e}_t, p_t)}(\overline{\gamma} - \underline{\gamma}) \approx g(\overline{\gamma} - \underline{\gamma}).$$

Hence by Theorem 1 the effort function may be written as $\hat{e}(p_{ti}, c_{ti}) = \overline{e} + bc_{ti} + g(\overline{\gamma} - \underline{\gamma})p_{ti}$. Second, we assume that the cost increment $c_{2i}(X_{1i}) - c_{1i} = c(X_{1i}) + \epsilon_{ci}$ has a common component and an idiosyncratic component, so that the impact of the first-period outcome on second-period cost is independent of the agent's cost in the first period. We can then write

$$D_i = b(C(1) - C(0)) + g(\overline{\gamma} - \underline{\gamma})(p_2(1, p_{1i}) - p_2(0, p_{1i})).$$

[7]Linear approximations are widely used in the applied literature and their validity has been extensively discussed: White [7], for example, has a good overview.

In a model without learning effects, this would deliver $D_i = D$ independent of i, but in our setting we also need to assume that there is no heterogeneity in p_{1i}. Assume for the moment that this is the case. Then we have

$$D_i = D = b(C(1) - C(0)) + g(\overline{\gamma} - \underline{\gamma})(p_2(1, p_1) - p_2(0, p_1)).$$

To estimate D the econometrician observes noisy signals of each agent's effort in the second period, denoted $y_{2i} = e_{2i} + \epsilon_{2i}$, where ϵ_{2i} is independent of other random variables and locations. The econometrician also observes either a noisy signal of effort in the first period $y_{1i} = e_{1i} + \epsilon_{1i}$ where ϵ_{1i} is independent of other random variables and location, or directly observes z_{1i} (but not both). In the former case the econometrician uses difference-in-differences to estimate the treatment effect, in the latter case they use a regression discontinuity analysis. As we now show in both cases under our assumptions, with enough data the econometrician can determine D with a high degree of reliability.

Take first the case where $y_{1i} = e_{1i} + \epsilon_{1i}$ is observed by the econometrician. In this case

$$y_{2i} - y_{1i} = e_{2i} - e_{1i} + \epsilon_{2i} - \epsilon_{1i} = g(\overline{\gamma} - \underline{\gamma})(p_2(X_{1i}) - p_1) + b\left(c(X_{1i}) + \epsilon_{ci}\right) + \epsilon_{2i} - \epsilon_{1i}.$$

Hence the subsample mean $\overline{d}(X_1)$ of $y_{2i} - y_{1i}$ for those i with a common value of X_1 is a consistent estimate of $d(X_1) = \hat{e}_0(p_2(X_1)) - \hat{e}_0(p_1) + b\left(c(X_1)\right)$, and since $D = d(1) - d(0)$, it is consistently estimated by $\overline{D} = \overline{d}(1) - \overline{d}(0)$. This is the difference-in-differences estimator.[8]

Second, take the case where the econometrician observes z_{1i}.. Here we add the assumption that c_{1i} is drawn from a continuous density function.[9] Assuming that the expectation exists. with enough data we can consistently estimate $E(y_2|z_1)$, which is

$$E(y_2|z_1) = \begin{cases} \overline{e} + g(\overline{\gamma} - \underline{\gamma})p_1(0) + bc_2(0) + E\left[c_1|z_1\right] & \text{if } z_1 < 0 \\ \overline{e} + g(\overline{\gamma} - \underline{\gamma})p_1(1) + bc_2(1) + E\left[c_1|z_1\right] & \text{if } z_1 > 0 \end{cases}.$$

It follows that for $z_1 > 0$ we have $E(y_2|z_1) - E(y_2| - z_1) = D + E\left[c_1|z_1\right] - E\left[c_1| - z_1\right]$. Since c_1 and z have a continuous densities and are mutually independent,

[8]Unmeasured changes are more likely to have the same effect on both groups if the groups have similar distributions of characteristics. Researchers typically use methods such as matching and synthetic controls to ensure this.

[9]Estimating the common effect with regression discontinuity analysis does not require that ϵ_{ci} be independent of other random variables, but the interpretation of the common effect as structural is problematic if ϵ_{ci} is correlated with the unobserved heterogeneity, as then a population with a different distribution of unobserved heterogeneity will respond differently to the treatment.

the joint distribution of c_1, z_1 is continuous so that for $z_1 \to 0$ we have $E[c_1|z_1] - E[c_1|-z_1] \to 0$. Hence for z_1 small the regression discontinuity $E(y_2|z_1) - E(y_2|-z_1)$ provides a good approximation to D.

Our warning about natural experiments can now be stated more precisely: The estimated treatment effect is

$$D = g(\overline{\gamma} - \underline{\gamma})(p_2(1) - p_2(0)) + b(c_2(1) - c_2(0)).$$

This confounds two effects, a learning effect $\Gamma \equiv g(\overline{\gamma} - \underline{\gamma})(p_2(1) - p_2(0))$ and a preference effect $b(c_2(1) - c_2(0))$ where $b < 0$. Moreover, Theorems 1 and 2 allow us to sign the learning effect and bound its magnitude. The next result is an immediate consequence of Theorems 1 and 2.

Corollary 1. *The learning effect satisfies*

$$0 < \Gamma \leq 2g(\overline{\gamma} - \underline{\gamma}) \frac{1 - \kappa_1}{\min\{F(-\overline{\gamma}), 1 - F(0)\}}.$$

The difference $c_2(1) - c_2(0)$, which describes a change in preferences and/or technology, may be structural, but the learning effect is not, because it depends on the information agents e have prior to treatment.. Suppose, for example, that empirical analysis established that there is a positive effect of the treatment, but that this was due to a strong learning effect and that $c_2(1) - c_2(0)$ is actually positive. If the treatment were then to be recommended and widely used, uncertainty about the effect of effort would diminish, so from Theorem 2 the learning effect would diminish as well, and the effect of the treatment would then be negative rather than positive.

The learning effect is always positive: success makes the agent think effort matters more and so raises effort. Hence the sign of D estimated by the econometrician matters. If it is negative, then it must be that $b(c_2(1) - c_2(0))$ is negative, and then D understates its magnitude, since there is an offsetting positive effect due to learning. On the other hand, if the estimated D is positive, we cannot be sure of the sign of $b(c_2(1) - c_2(0))$, since the positive value of D may be due to the positive effect of learning offsetting a smaller negative structural effect. In other words, negative estimated values of D are robust to learning and so are more useful for policy purposes than positive ones are.

Corollary 1 shows that the learning effect is small if $\overline{\gamma}$ and $\underline{\gamma}$ are close, or if the agent believes that one of the two states is very likely. In either of these cases difference-in-differences or regression discontinuity analysis will, to a good approximation, yield an estimate of $b(c_2(1) - c_2(0))$ which can be regarded as structural. Moreover, when this is the case, the estimates are also robust to heterogeneity in first-period beliefs.On the other hand, if the belief effect is large the interpretation of the treatment effect is problematic even if beliefs are homogeneous.

The learning effect will also be small if the agents observe z_1, as they may if the underlying score variable is a test score or the vote differential in an election. In this case success or failure gives agents no additional information about the efficacy of effort beyond that contained in z_1, so their posterior is $p_2(z_1)$ independent of X_1. As $p_2(z_1)$ is continuous in z_1, this implies that near the threshold where $z_1 = 0$ there is no learning effect, and indeed the regression discontinuity will be equal to $b(c_2(1) - c_2(0))$. If z_1 is observed with error, but the error is small, then the belief effect will also be small.

5. Applications

Here we examine two applications. In each case we examine studies that have received substantial attention in the form of publication in top journals, citations, and discussion in the media.

5.1. State Violence

The effect of state violence on insurgencies has received considerable attention in the political economy literature. The setting is one of civil strife where a central government faces an insurgency and bombs or shells agents who correspond to villages. Here effort e corresponds to resisting the insurgents and $X_1 = 0$ corresponds to being bombed. *A priori* there are competing theoretical predictions of the effect of the bombing "treatment." One theory is that bombing angers villagers and makes them less sympathetic to the government and more sympathetic to the insurgents, so bombing makes villagers less inclined to resist. This corresponds to a change in preferences in which $c_2(0) > c_2(1)$, so bombing lowers the marginal utility of effort. A competing theory is that bombing destroys infrastructure and makes the village less attractive to insurgents, so it becomes easier for villagers to resist the insurgents. This corresponds to a reduction in the cost of resistance, so $c_2(0) < c_2(1)$.

Dell and Querubin [2] uses a novel data set based on the indices used by the US Air Force to decide which hamlets to bomb during the Vietnam War. Villages were assigned scores that researchers observed s (with some difficulty) but villagers did. Villages were then put into a small number of categories based on their scores and bombed accordingly.[10] Hence if we compare villages very near the threshold

[10] Our simple model has only a single threshold; the next section briefly discusses the extension to multiple thresholds.

we should should be able to estimate the treatment effect of bombing by regression discontinuity analysis, and Dell and Querubin [2] do this. They find that "[b]ombing increased the military and political activities of the communist insurgency, weakened local government administration and lowered non-communist civic engagement."[11]

Our analysis indicates that the negative treatment effect of bombing on compliance that Dell and Querubin [2] identified was due to learning, and that the structural effect on preferences was either zero or positive. That is, the villages that were bombed may have concluded effort did not matter very much and so reduced their effort. This always lead to a negative estimated effect if the structural effect is zero, and can lead to negative estimates despite a positive structural effect if the villagers were sufficiently uncertain about how much effort matters.

We think there is evidence that villagers were indeed fairly uncertain about the links between their effort and their probability of being bombed. It is true that t, the U.S. dropped leaflets to advertise that effort matters (Dell and Querubin [2]), which s might well lead villagers to believe that effort matters, the leaflets did not answer the question of the size of its effect. Moreover, while villagers might have had some information about whether other villages were bombed, they seem less likely to have had information about how much effort those villages made, so their own experience of being bombed or not would still have been an important source of information about the strength of incentives. We conclude that we cannot tell from Dell and Querubin [2] whether the induced change was the result of a structural change in preferences or a more ephemeral effect of learning. Ir To assess the structural effect of state violence, researchers need a setting where there was a high degree of certainty about the efficacy of effort. This appears to have been the case in the study of Lyall [5] using Chechen data. Here Russian military doctrine called for random and unpredictable shelling independent of effort. It appears that not only was this doctrine adhered to, but was enhanced by the fact that the targeting of villages was largely performed by soldiers who had been drinking heavily. Of course the mere fact that targeting was indiscriminate does not indicate that villagers knew this, but Lyall [5] gives evidence that this was the case, ranging from the fact that the drunken behavior of Russian soldiers was well known to villagers to the fact that they made formal complaints about being targeted for no reason. In other words, in the Lyall [5] study there is reason to believe that learning effects were small. It is striking then that Lyall [5] reaches the opposite conclusion from Dell and Querubin [2]: Using a

[11] They also compare the punishments used in U.S. Army administered areas to the rewards used in U.S. Marine Corp areas. As this is a different question we do not comment on that analysis here.

difference in difference method[12] Lyall [5] finds that after shelling insurgent activity decreased.

There are many differences between Chechnya and Vietnam, including their prevailing religious and cultural norms, but it appears to be the general view in the state-violence literature that lessons learned from one conflict are applicable to another. If this is the case, then a possible conclusion is that bombing and shelling does work in the sense of increasing effort, but this effect may be offset in the short run when the people being bombed are uncertain how much their efforts matter. Notice that this conclusion reconciles the view that state violence works- which is presumably what motivated the bombing campaign-with the empirical findings that it does not. .

5.2. Incumbency

It is widely believed that incumbents have an electoral advantage. To fit this into our framework, suppose that the agent is a candidate for elective office. The outcome $X_1 = 1$ corresponds to winning the election, and $X_1 = 0$ to losing it. Effort corresponds to campaign effort. Here the treatment corresponds to $X_1 = 1$ and an incumbency effect means that winning an election gives the incumbent an advantage in subsequent elections, perhaps because incumbents receive free publicity by virtue of their office, or because they can do favors for their constituents. All of this lowers the cost of effort provision $c_2(1) < c_2(0)$, so should increase effort and chances of success in subsequent elections. Note however that there could also be an effect in the opposite direction: For example, it might be that personal campaigning is important and an incumbent politician who lives and works in Washington D.C. might it costly to campaign in her local district.

This issue was studied by Lee [4] using a regression discontinuity analysis. Here the latent variable z_1 is the vote differential and naturally those who just lose an election are not so different from those who just win. Lee [4] finds a substantial discontinuity indicating that just winning greatly increases the chances of future success as against just losing.

As in Dell and Querubin [2], effort provision here is endogenous and depends upon beliefs about the efficacy of effort in turning out votes. Moreover, this goes

[12] Strictly speaking if the shelling was indiscriminate difference in differences was not needed, as the treated and untreated villages should have been the same. Lyall [5] is conservative in this respect, although in fact his data shows very little difference in insurgent activity (effort) in the treated and untreated villages prior to shelling. This is consistent with the idea that the shelling was indeed indiscriminate.

in the "wrong" direction for Lee: The learning effect means that winners would provide more effort and losers less - not a true incumbency effect through the cost function. However,: unlike Dell and Querubin [2] there is reason here to believe that the learning effect is small. First, political parties have a lot of data about past elections, so probably have a pretty good idea about the efficacy of effort. Second, even if this is not the case, the latent variable z_1 is directly observed, so in our the actual success or failure conveys no additional information and we predict so the learning effect will not be present.[13]

6. Discussion

6.1. Incentives and Beliefs

The dependence of treatment on effort creates an incentive effect that induces increased effort. For example, admission to an exclusive college as a reward may induces high school students to study more if they extrapolate their college application experience to the rewards to effort in college. Here the benefit of treatment is not simply that the treated group behaves differently than the non-treated group, but also that it provides incentives for effort both before and after treatment. This implies policy evaluations may have to weigh increased effort due to *ex-ante* incentives against decreased effort *ex-pos*t due to failure. As we shall show, it matters in this assessment whether the treatment effect is due to a change in preferences or to learning.

Specifically, if treatment reduces *ex post* effort, then it seems as if there is a tradeoff between the reduced effort due to failure and the increased effort *ex-ante* due to incentives. However, if the reduction in *ex post* effort is due to learning, then while it is true that agents who fail lower their effort level, they still provide a higher level of effort than they would have if they believed that effort had little effect. A good example of this is the question of how and whether state violence "works." Answering the narrow question of whether villages that are bombed as part of a bombing campaign comply more or less than villages that are not bombed does not answer the larger question of whether a bombing campaign may be desirable because the threat of bombing induces compliance.

The agents' beliefs about the link between their effort and their success play two roles here. First, the agents' initial beliefs about the efficacy of effort determines

[13] Note that direct observation of the underlying index would not have this consequence if the threshold itself was uncertain, as then success or failure would contain additional information. This doesn't seem relevant in the elections Lee [4]studies, as the candidates presumably knew the rules.

their initial effort provision, and second their updating these beliefs leads the realized treatments to have a learning effect. We do not expect these beliefs to be arbitrary, but to have some basis in how effective effort is. For example, if effort does not matter we expect that people generally understand that it does not matter much. This was our analysis of the Lyall [5] application.

In the case of a policy designed to bring about changes in behavior, it well understood that keeping the policy secret is a bad idea. This is why for example many houses with burglar alarms make that fact known, and why visible monitors of traffic speed are used to get drivers to slow down. [14] However, simply announcing a deterrence policy may not be sufficient, because the policy maker may face a credibility problem: It would like to persuade the agents that effort matters a lot but would prefer not to pay the cost associated with providing incentives. For the incentive policy to be effective, the policy maker must therefore clearly communicate its commitment to the policy, and then stick to it at least most of the time.

Our analysis here shows a second undesirable consequence of secrecy: Uncertainty about how much effort matters that creates the undesirable learning effect. This effect can be mitigated by providing detailed and credible information about how punishment depends upon effort - or by simply revealing the score z_1.

6.2. Multiple Treatment Thresholds

In examining regression discontinuities we have assumed that there is a single treatment and a single threshold. There might in fact be several: For example, if z_1 represents a normalized test score, then a positive value might mean admission to an exclusive college, while a score between -1 and 0 might mean admission to a less exclusive college, and a score less than -1 might mean not being admitted to any college.

The presence of multiple thresholds is potentially complicated as it can lead to non-convexities in the agent's objective function. Here we assume that agents fall into several groups, perhaps corresponding to "good students" (low cost of effort) and "poor students" (high cost of effort). The idea is that for the "poor students" the relevant threshold is the one between no college and college, while for the good student the relevant threshold is between a more or less exclusive college. In this case we can apply our analysis separately to each group, assuming each faces a single threshold.

[14] In the words of Dr. Strangelove: "The whole point of the doomsday machine...is lost if you keep it a secret!

The key point is that multiple thresholds enable us to observe the treatment effect for different effort levels. That is, the group of poor students produces lower effort than the group of good students. Moreover, from Theorem 2, the strength of the learning effect is greater at higher effort levels where the signal is more revealing. Hence if learning is important we should see a stronger treatment effect for the higher effort groups.[15]

6.3. Multi-dimensional Learning

Our model of learning is designed to capture the idea that success signals that effort matters and failure signals that it does not. This follows from the fact that agents are uncertain about the rate at which effort lowers the chances of failure. They are, however, certain about the probability of success when no effort is made. Consider the case where the probability of a bad outcome is of the form $qF(-\gamma e)$ where q is uncertain but γ is known. Then a success is evidence for lower q. However, lower q means that effort will have less effect on the outcome, so the incentives for effort are weaker. In this setting success will lead to a reduction in effort, which is the opposite of what happens when agents learn aboutγ. In practice both effects will be present, and their relative strength depends both on the degree of uncertainty about the two parameters q, γ and on the level of effort provided. When little effort is provided little information about γ is generated by success or failure: Indeed, in the extreme case of where no effort is made, success or failure provides information about q but not about γ. On the other hand, our analysis shows that the larger the effort the more information is provided about γ. It is natural to conjecture that under appropriate technical conditions there should be an effort cutoff e^* below which success will reduce effort, and above which it raises effort. The more information available about q the lower will be the threshold e^*.

In general we expect that agents have more precise and accurate beliefs about the overall chance of success than about how success is influenced by effort. In particular, this seems plausible in the Dell and Querubin [2] and Lee [4] examples discussed above, especially in Dell and Querubin [2] where most of the data is for high-effort villages. However, this may not be true in other contexts. For example, for a homeowner, "failure" might be a burglary. A homeowner with a low cost of protecting against burglars will engage in a high level of effort. If their home is

[15] In Dell and Querubin [2] there are in fact multiple thresholds. In their data the lower effort hamlets do appear to have a smaller treatment effect, but because there is relatively little data on low -effort hamlets that were not bombed, the difference between groups is not statistically significant. And of course there are other possible explanations for this, for example the utility loss of being bombed might have been lower in low effort villages.

broken into they will conclude that effort does not matter very much, and reduce their effort. By contrast a homeowner with a high cost of protecting against burglary will engage in a low level of effort. Thus break-ins will lead them to infer that burglary is more likely than they had suspected and increase their level of effort.

The possibility of a threshold below which success signals effort does not matter raises some additional issues relevant for empirical work. On the positive side, if the sample is split more or less equally between those above and below the effort threshold e^*, then the positive and negative effects may more or less cancel out so that the learning effect is small. Conversely, if the sample is largely below the threshold then the learning effect is reversed, so that if we are uncertain whether agents are more likely to be above or below the threshold, we cannot even be certain of which direction the learning effect goes, so the estimated threshold effect cannot be taken as a bound on the true structural effect.

7. Conclusion

The use of natural experiments in empirical work has greatly contributed to our understanding of economic phenomena by directing researchers to obtain more revealing data and better instruments, along with improved and clever techniques for taking advantage of that better data. Here we have pointed out, as have others, that this is a complement for theory, not a substitute. Randomness and explanatory power are necessary but not sufficient for the identification of a structural parameter, because the interpretation of empirical analyses typically requires either an explicit model or an implicit one.[16] Specifically, when unobserved heterogeneity involves effort, the "treatment" may provide agents with useful information that influences their behavior. As we have shown, this can lead to quite different policy recommendations than when learning effect is absent. In addition we have suggested cases where the learning effect is likely to be negligible, for example when using regression discontinuity analysis to analyze elections, the success or failure of a candidate do not add information to the vote differential.

[1] Joshua David. Angrist and Jorn-Steffen. Pischke. *Mostly harmless econometrics : an empiricist's companion.* Princeton University Press, 2009. ISBN 1400829828. URL https://books.google.com/books/about/Mostly_Harmless_Econometrics.html?id=ztXL21Xd

[16]See, for example, Heckman and Smith [3] and Rosenzweig and Wolpin [6].

[2] Melissa Dell and Pablo Querubin. Nation Building Through Foreign Intervention: Evidence From Discontinuities in Military Strategies. *Quarterly Journal of Economics*, 133(2):701–764, 2018. ISSN 0033-5533. doi: 10.1093/qje/qjx037.

[3] James J Heckman and Jeffrey A Smith. Assessing the Case for Social Experiments. *Journal of Economic Perspectives*, 9(2): 85–110, 1995. ISSN 0895-3309. doi: 10.1257/jep.9.2.85. URL http://pubs.aeaweb.org/doi/10.1257/jep.9.2.85.

[4] David S. Lee. Randomized experiments from non-random selection in U.S. House elections. *Journal of Econometrics*, 142(2):675–697, 2008. ISSN 03044076. doi: 10.1016/j.jeconom.2007.05.004.

[5] Jason Lyall. Does Indiscriminate Violence Incite Insurgent Attacks? Evidence from Chechnya. *Journal of Conflict Resolution*, 53(2):331–362, 2009. doi: 10.1177/0022002708330881.

[6] Mark R Rosenzweig and Kenneth I Wolpin. Natural "Natural Experiments" in Economics. *Journal of Economic Literature*, 38(4):827–874, 2000. ISSN 0022-0515. doi: 10.1257/jel.38.4.827. URL http://pubs.aeaweb.org/doi/10.1257/jel.38.4.827.

[7] Halbert White. Using least squares to approximate unknown regression functions. *International Economic Review1*, 21(1):149–170, 1980. URL https://www.jstor.org/stable/2526245.