

Learning in Games and the Interpretation of Natural Experiments¹²

Drew Fudenberg³ and David K. Levine⁴

Abstract

We show that the treatment effect estimated by standard methods such as regression discontinuity analysis or difference-in-differences may contain a transient “learning effect” that is entangled with the long term effect of the treatment. This learning effect occurs when the variable of interest is the agents’ efforts, when treatment and control correspond to success or failure: Success or failure gives agents information about how much their effort matters, and consequently changes the amount of effort they provide after treatment. We examine the impact of the learning effect and when it is likely to be substantial.

¹We thank Chris Ackerman, Daron Acemoglu, Josh Angrist, Daniel Clark, Amy Finkelstein, Michael Griffiths, Kevin He, Andrea Ichino, Andrea Mattozzi, Parag Pathak, Demian Pouzo, Leeat Yariv, and the WUSTL Economics workshop for helpful comments and conversations, and two referees for very helpful reports. We are grateful to NSF grant SES 1643517 and the EUI Research Council for financial support.

²This version: 3/17/2020. First version: 01/02/2019

³Department of Economics, MIT; drew.fudenberg@gmail.com

⁴Department of Economics and RSCAS, EUI, and Department of Economics, WUSTL; david@dklevine.com

Natural experiments are widely used in economics to estimate treatment effects in non-experimental settings. As is well known, randomness and explanatory power are necessary but not sufficient for the identification of a structural parameter, because a variety of factors can impede the straightforward mapping from reduced form estimates to stable parameters of interest.⁵ This paper points out another issue to consider in the construction and interpretation of these experiments, and discusses the types of situations where this issue is more likely to be important: When selection into the treatment depends upon the subjects' effort, the treatment effect may depend on their beliefs, and as in the Lucas [19] critique of the Phillips curve, these beliefs may change as the subjects gain experience. Thus causal interpretations of the treatment effect may not be as convincing as might have been supposed.

We focus on natural experiments where the "treatment" corresponds to "success" or "failure," and the subjects choose effort levels both before and after the treatment. 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 subjects observe whether they succeed or fail, but do not observe the underlying "score" variable that is used to determine their treatment, so that success indeed signals that effort is likely to be effective, and failure signals that it probably has less effect. As we explain, this effect of learning on second-period effort 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 treatment effect 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.

Chassang et al. [3] makes a related point: a treatment may have both a direct effect, for example, bombing a village may destroy infrastructure useful to terrorists, and an indirect effect through effort decisions, for example, bombing a village may antagonize villagers so that they work less hard to prevent terrorist activity, but if the subjects have heterogeneous initial beliefs about the efficacy of effort, an RCT

⁵See for example Sims [24], Heckman [13] and Rosenzweig and Wolpin [22].

cannot disentangle these two effects. We abstract from this by assuming that there is only an indirect effect, or what amounts to the same thing, that the econometrician is interested only in the combined direct and indirect effect. Hence our critique, unlike Chassang et al. [3] and related work such as Chassang et al. [4], Chemla and Hennessy [5], and Philipson and DeSimone [21] does not apply to randomized control trials. By contrast we identify a new learning effect that poses a problem for natural experiments even if initial beliefs are homogeneous. In a similar vein, we assume that the treatment effect of interest is homogeneous, so that the local average treatment effect is global, which avoids the potential confounds in regression-discontinuity estimates that was pointed out by Deaton [6]. We show that even absent this confound learning effects create a confound of their own when the probability of being assigned the treatment is endogenous, and depends both on the agent’s effort and on an unknown state of the world

From a policy perspective 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.

We use a very simple and stylized model to develop a precise characterization of the confounds that can arise from neglecting learning effects in two leading methods for analyzing natural experiments, namely regression discontinuity and difference-in-differences.⁶ In our model, both of these methods work perfectly in the absence of learning, and indeed both the average treatment effect and the local average treatment effect reduce to a simple “treatment effect” that is the same for all subjects. However, the estimates from these methods combine a persistent preference effect and a more transient learning effect. We analyze the difference between the overall treatment effect and its persistent component. We show that regression discontinuity consistently estimates a local average treatment effect, and that when all agents start with the same prior beliefs the difference-in-differences estimator gives an intermediate value between the average treatment effects for the treated and untreated. In both cases, we give conditions under which the learning effect will be relatively small. We also distinguish between a direct learning effect on effort and possible long-term indirect consequences due to investment.

We apply our analysis to three well-known natural experiments, those of Dell and Querubin [7], Lee [16], and Hoekstra [15]. Dell and Querubin [7] uses a novel data set based on the indices constructed by the US Air Force to decide which hamlets to

⁶For definitions of these procedures see Angrist and Pischke [2] for example.

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 might suggest that the bombing campaign was counterproductive, but we show that this conclusion does not follow. Intuitively, the idea is that the threat of bombing may induce compliance, while actually being bombed may convince villagers that compliance is pointless as they will be bombed anyway. Thus even though the assignment is locally randomized, and the estimated average treatment effect is consistently estimated, the implications of the estimated effects depend on the model in which they are evaluated. Additional insight can be derived from Lyall [20], which studies the indiscriminate shelling of Chechen villages by the Russians, where this learning effect is less likely to be present, and reaches the opposite conclusion from Dell and Querubin [7]. This contrary finding raises the possibility that in Dell and Querubin [7] the preference effect may have been positive, while the learning effect was negative. Because the econometrician only observes the combined effect of these two forces, the coefficient estimates must be interpreted with care.

Lee [16] 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.⁷ Lee 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, but we argue that the learning effect is likely to be small in this setting. In particular the candidates typically observe the vote differential, and in this case success and failure provide no additional information about the benefits of effort.

Hoekstra [15] studies the effect of admission to a flagship college on subsequent earnings, and finds a substantial impact based on a regression discontinuity analysis. The time lag between the treatment and its measured effect is long enough that it is hard to attribute this to a direct effect of learning. However, this could be due to an indirect learning effect: those who just miss admission may conclude that effort makes little difference and invest little in human capital during their college years, resulting in lower earnings. Even if subsequently the “losers” discover that effort matters, it may be too late to make up the human capital deficiency.

Section 1 of the paper lays out and analyzes our theoretical framework, in which

⁷Empirical studies had previously established that incumbents have an electoral advantage: see for example Gelman and King [12]. But this leaves open the possibility that incumbents are incumbents simply because they are better politicians, and not because of any intrinsic advantage of incumbency. Lee uses a regression discontinuity analysis to address the issue of causality.

the subjects are agents who need to decide how much effort to exert, and are uncertain about how effort changes the probability of success. Section 2 examines the econometric issues, and computes the implications of our model for the estimated treatment effects in regression discontinuity and difference-in-differences 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 3 then applies our findings to two empirical settings: the study of state violence and incumbency. Section 4 discusses broader issues including the robustness of our results.

1. Theoretical Framework

1.1. The Model

We study a two period model $t = 1, 2$ with many agents i . In period 1 each agent is assigned to a treatment $D_{i1} \in \{0, 1\}$. The treatment is determined by an assignment variable x_{i1} , where the agent is treated when the assignment variable is positive, that is $D_{i1} = 1$ if and only if $x_{i1} > 0$. In the applications we consider, x_{i1} might represent the difference between a score and a threshold.⁸ In period 2 each agent observes their treatment D_{i1} and their outcome y_{i2} , but not their assignment variable x_{i1} . They do not observe the outcomes, assignments, or treatments of others.⁹

Our setting is a natural experiment in which the situation of the agent repeats each period: Each period $t = 1, 2$, agent i determines a level of effort $e_{it} \in [0, 1]$. Effort then influences the outcome variable $y_{it} = e_{it} + \epsilon_{it}$ where the ϵ_{it} have zero mean and are independent of all other shocks. The assignment x_{it} and “treatment” D_{it} also recur each period; for example, if the treatment is the outcome of an election, there will be a subsequent election. Effort also influences the assignment variable according to the formula $x_{it} = \gamma e_{it} + \omega_{it}$, where $1 \geq \gamma \geq 0$ measures the efficacy of effort and ω_{it} is an iid random shock with cdf F .¹⁰

Effort in each period has benefits and costs. The key econometric problem is to infer how these costs and benefits are changed by the treatment: For example, does winning an election create an incumbency advantage? We divide the utility of effort e_{it} into three additive components. First the “treatment” from the perspective of the

⁸We discuss the decomposition of x_{it} in Section 4.4.

⁹As we discuss in Section 4 below, this is a natural assumption in the settings where we find that learning effects may be important. If agents observe the outcomes of others, this provides an additional source of information to the extent that they can observe or infer the actions of others, and so weakens the learning effect based on the agents’ direct observations of their own outcomes.

¹⁰Note that we do not expect the random variable ω_{it} to have zero mean.

agent is utility relevant. We take $D_{it} = 0$ to represent failure (losing the election) and $D_{it} = 1$ to represent success (winning the election), and assume that there is an additive bonus of 1 for achieving success. Second, there is an expected benefit of effort $U(e_{it})$.¹¹ We assume that the functional form of U is the same for all agents, and incorporate all individual-specific differences in the cost component. Exerting effort has constant marginal cost of $c_{i1} > 0$ in period 1, and constant marginal cost of $c_{i2}(D_{i1})$ in the second period. In this formulation, $c_{i2} - c_{i1}$ captures all of the changes in preferences and productivity brought about by the treatment D_{i1} . We assume that the econometrician's goal is to learn $c_{i2}(D_{i1}) - c_{i1}$, which captures the effect of the treatment on the underlying cost.

The agents know the distribution of ω_{it} but are uncertain about the effectiveness of effort, which is measured by γ . For simplicity we assume that agents contemplate only two possible levels of γ , namely $\bar{\gamma}$ and $\underline{\gamma}$ with $1 \geq \bar{\gamma} > \underline{\gamma} \geq 0$. Agent i assigns prior probability $1 > p_{i1} > 0$ to the state of the world $\gamma = \bar{\gamma}$ where effort matters more.

Our assumptions imply that the only information each agent i observes about the effectiveness of effort is their own treatment D_{i1} . Let $p_{i2}(D_{i1})$ denote the agent's posterior after observing D_{i1} , as determined by Bayes law.¹² To simplify the analysis, we assume that agents are myopic and choose e_{it} solely to maximize utility in period t . This means that in period 2 they use their information from the treatment in period 1, but do not alter their choice of e_{i1} to have a more informative signal about how much effort matters. There are two reasons that this is plausible in the settings we examine. First, the informational benefit of investing in greater effort is small, and second, the benefits of that information are relatively long-run in nature, hence highly discounted. For example, we do not think that high school students' choice of scholarly effort is likely to be influenced by the consideration that additional effort will better reveal the benefits of effort some decades in the future. Moreover, as we explain in Section 4.5, our results for regression discontinuities apply immediately to the case of non-myopic agents, though the results for difference-in-differences become less sharp.

We now make a strong functional form assumption that will eliminate many complications in the econometric implementation of the model, and will let us focus

¹¹Note that this is consistent with a formulation in which utility depends on the outcome y_{it} and $U(e_{it})$ is expected utility conditional on effort e_{it} .

¹²Note that when $e_{1i} = 0$ 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.

attention on the econometric implications of learning by the agents.

Assumption. (*Linear-Quadratic*)

- a. $U(e_{it}) = u(e_{it} - e_{it}^2/2)$ for a scalar constant $u > 0$ and $e_{it} \in [0, 1]$.
- b. The support of F includes $[-1, 0]$, and in this range $F(\omega) = F_0 + f\omega$, with $F_0 > f > 0$.
- c. $\underline{\gamma}f + u > c_{it} > \bar{\gamma}f$.

We will show that these assumptions imply that the treatment effect is the same for all agents, and is the same as the average treatment effect estimated by difference-in-differences estimation and the local average treatment effect estimated by regression discontinuity analysis. The linear-quadratic functional forms may be viewed as a convenient approximation and will lead to effort provision that is linear in cost and beliefs, a property that is often used in applications.¹³ The final part of the assumption is that cost is of intermediate size: this will ensure the existence of an interior solution to the optimal effort problem.

Our ultimate goal is to analyze what an econometrician can learn from observing this system about the effect of treatment (D_{i1}) on second-period effort e_{i2} and in particular what part of the observed effect is driven by changes in cost, which are likely to be persistent, and what part of the observed effect is due to learning and likely to be more transient. We discuss this further when we introduce the econometrician formally in Section 2.

1.2. Learning and Optimal Effort

Agent i succeeds in period t if $\gamma e_{it} + \omega_{it} > 0$, which has probability $1 - F(-\gamma e_{it})$. Thus the objective function for agent i in period t is

$$v(e_{it} | p_{it}, c_{it}) = p_{it}(1 - F(-\bar{\gamma}e_{it})) + (1 - p_{it})(1 - F(-\underline{\gamma}e_{it})) + U(e_{it}) - c_{it}e_{it}.$$

We will consider three applications. One is a setting of civil strife, where a central government faces an insurgency and bombs or shells villages. Here effort corresponds to resisting the insurgents, and “success” corresponds to not being bombed, which is $D_{it} = 1$. 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

¹³If we let $F(\omega)$ be quadratic then effort would not be linear in beliefs.

in preferences in which $c_{i2}(0) > c_{i1} > c_{i2}(1)$, so that bombing lowers the marginal utility of effort. Another theory is that bombing destroys infrastructure and makes the village less attractive to insurgents, which makes it easier for villagers to resist. This corresponds to a reduction in the cost of resistance so that $c_{i2}(0) < c_{i1} < c_{i2}(1)$.

In the second application the setting is an election, and the agent is a candidate. Here treatment $D_{i1} = 1$ is winning the first-period election and $D_{i1} = 0$ is losing it. In this case one theory is that winning the election gives the winner, the incumbent, an advantage in the subsequent election. This is the same as lowering the cost of effort and corresponds to $c_{i2}(1) < c_{i2}(0)$. 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.

The third application is college admission, and the agent is an applicant. Here the treatment $D_{i1} = 1$ is being admitted to a prestige school. One theory is that admission to the prestige school gives the student a higher income y_{i2} due to better education, better connections and so forth. In our model this corresponds to lowering the cost of effort, i.e. $c_{i2}(1) < c_{i2}(0)$. Alternatively it might make little difference whether a student attends a prestige school or a slightly less prestigious school.

For given values of c_{it} and p_{it} agent i faces a simple static optimization problem. In our linear quadratic model this has a unique closed-form solution $\hat{e}(p_{it}, c_{it})$. Our first result gives the key properties of that solution:

Theorem 1. *Under the linear quadratic assumption the unique solution of agent i 's problem in period t is linear in i 's cost and i 's belief, and is given by*

$$\hat{e}_{it} = \Gamma_0 - \Gamma_c c_{it} + \Gamma_\ell p_{it}$$

where $\Gamma_0 = 1 + \underline{\gamma}f/u$, $\Gamma_c = (1/u)$, and $\Gamma_\ell = (f/u)(\bar{\gamma} - \underline{\gamma})$.

Proof. The derivative of the objective function is

$$p_{it} [\bar{\gamma}f - \underline{\gamma}f] + \underline{\gamma}f + u - ue_{it} - c_{it}.$$

The second derivative is $-u < 0$ and by assumption this is positive at the lower boundary $e_{it} = 0$ and negative at the upper boundary $e_{it} = 1$ so the solution is the interior one given in the theorem. \square

Our second main result is that the posterior probability that $\gamma = \bar{\gamma}$ is higher after success, and that this effect is weaker when the agent's prior is closer to 0 or 1 as measured by $\kappa_{i1} = |1 - 2p_{i1}|$. We also show that the way the agents update their

beliefs depends on their effort, and that higher levels of e_1 are more informative. These facts will have important consequences when we analyze the sensitivity of the estimated treatment effects to the confounding effects of the agent's inference process.

Theorem 2. *Under the linear quadratic assumption*

1. $p_{i2}(1) > p_{i1} > p_{i2}(0)$.
2. $p_{i2}(0)$ is strictly decreasing in e_{i1} , and $p_{i2}(1)$ is strictly increasing in e_{i1} .
3. $p_{i2}(D_{i1})$ is strictly increasing in F_0
4. The change in posterior satisfies

$$|p_{i2}(D_{i1}) - p_{i1}| \leq \frac{1}{\min\{F(-\bar{\gamma}), 1 - F(0)\}}(1 - \kappa_{i1}).$$

Result (1) says that success ($D_{i1} = 1$) is evidence that effort matters more. This reflects the fact that for any given effort level success is more likely if effort matters more.

Result (2) says that the signal is more informative when there is higher effort. This comes from the assumed multiplicative interaction between effort and the state of the world.

To understand result (3), suppose the threshold for success is not 0 but χ , so that failure occurs when $\gamma e_{it} + \omega_{it} \leq \chi$. Then the probability of failure is $F(\chi - \gamma e_{it}) = F_0 + f\chi - \gamma e_{it}$, and we see that an increase in the threshold increases the probability of failure in the same way as an increase in F_0 . A higher probability of failure means that a failure has less of an effect on the posterior, and so failure lowers the posterior less in the case of failure and raises it more in the case of success.

The final result says that greater certainty results in less change in the posterior.

Proof. Let $\Pr(D_{i1} | \gamma)$ denote the conditional probability of the treatment D_{i1} given γ . We can write Bayes law as

$$\begin{aligned} p_{i2}(D_{i1}) &= \left(\frac{\Pr(D_{i1}|\bar{\gamma})}{p_{i1} \Pr(D_{i1}|\bar{\gamma}) + (1 - p_{i1}) \Pr(D_{i1}|\underline{\gamma})} \right) p_{i1} \\ &= \left(\frac{1}{p_{i1} + (1 - p_{i1}) \Pr(D_{i1}|\underline{\gamma}) / \Pr(D_{i1}|\bar{\gamma})} \right) p_{i1} \end{aligned} \tag{1}$$

which is increasing in the likelihood ratio $L(D_{i1}) \equiv \Pr(D_{i1}|\bar{\gamma}) / \Pr(D_{i1}|\underline{\gamma})$. For $D_{i1} = 0$ we have

$$L(0) = \frac{\Pr(0|\bar{\gamma})}{\Pr(0|\underline{\gamma})} = \frac{F_0 - \bar{\gamma}e_{i1}}{F_0 - \underline{\gamma}e_{i1}} < 1,$$

and for $D_{i1} = 1$ we have

$$L(1) = \frac{\Pr(1|\bar{\gamma})}{\Pr(1|\underline{\gamma})} = \frac{1 - F_0 + \bar{\gamma}e_{i1}}{1 - F_0 + \underline{\gamma}e_{i1}} > 1,$$

from which it follows that $p_{i2}(1) > p_{i1} > p_{i2}(0)$, which proves (1).

To prove (2), observe that $L(0)$ is strictly decreasing in e_{i1} , and $L(1)$ is strictly increasing in e_{i1} , so $p_{i2}(D_{i1})$ has the same properties.

To prove (3), observe that $L(0)$ and $L(1)$ are both strictly increasing in F_0 .

To prove the final claim in the theorem, note that from Bayes law in equation 1

$$\begin{aligned} \left| \frac{p_{i2}(D_{i1}) - p_{i1}}{p_{i1}} \right| &= \left| \frac{(1 - p_{i1})(\Pr(D_{i1}|\bar{\gamma}) - \Pr(D_{i1}|\underline{\gamma}))}{p_{i1} \Pr(D_{i1}|\bar{\gamma}) + (1 - p_{i1}) \Pr(D_{i1}|\underline{\gamma})} \right| \\ &\leq \left| \frac{1}{\min\{\Pr(D_{i1}|\bar{\gamma}), \Pr(D_{i1}|\underline{\gamma})\}} \right|. \end{aligned}$$

The result then follows from minimizing the probabilities in the denominator over \hat{e}_{i1} .

By reversing the role of $\underline{\gamma}$ and $\bar{\gamma}$ the bound also holds for $1 - p_{i1}$. Since $\min\{p_{i1}, 1 - p_{i1}\} = 1 - |1 - 2p_{i1}|$ the result follows. \square

2. The Econometrician's Problem

Now we turn to the problem faced by an econometrician who has tools ke to estimate the effect of D_{i1} on \hat{e}_{i2} . As we shall see this estimate will combine two effects, a preference effect and a learning effect, that have very different implications for both persistence and for policy. Consequently the econometrician would like to isolate these two effects, but as we show this separation is typically not possible.

Once the econometrician has determined what they want to learn, the next step is to develop estimators that are consistent for particular population moments. Hence averages over the population, integrating out the agent labels i , are an important part of the analysis. To define these averages we postulate that agents i with costs c_{i1} and priors p_{i1} are drawn from an underlying probability distribution that we assume has a continuous density on an open set.¹⁴

The distribution on costs and priors, combined with the decision rule for effort $\hat{e}_{i1} = \Gamma_0 - \Gamma_c c_{1t} + \Gamma_\ell p_{i1}$, the definition of the assignment variable $x_{i1} = \gamma \hat{e}_{i1} + \omega_{i1}$,

¹⁴Occasionally it will also be useful to consider the case of common cost or common prior in which one of c_{i1} or p_{i1} , but not both, is constant.

outcome variable $y_{it} = e_{it} + \epsilon_{it}$ and Bayes Law induce a joint distribution over the assignment variables, priors p_{i1} , outcomes y_{it} , treatments D_{it} , and the posteriors $p_{i2}(0), p_{i2}(1)$. In this section we will take conditional expectations with respect to this joint distribution. For example, if \mathcal{X} is a cylinder set event corresponding to realizations of the treatment variable, we will write $\mathbb{E}[p_{i2}(0) \mid \mathcal{X}]$, where this expectation does not depend on the agent label i , which has been integrated out.

To eliminate many complications in the econometric implementation of the model, and let us focus on the implication of learning, we assume that the impact of first-period treatment on second-period costs is the same for all of the agents.

Assumption. (*Limited Heterogeneity*) $c_{i2}(D_{i1}) = c_{1i} + \mathcal{C}(D_{i1}) + \epsilon_i$, where the ϵ_i are *i.i.d.* shocks with zero mean, and $\mathcal{C}(D_{i1})$ is a common effect of the treatment on preferences.

Treatment, Preference, and Learning Effects

The *treatment effect* of D_{i1} on \hat{e}_{i2} is the expected difference in second-period effort of two agents with identical period-1 characteristics, one of whom is treated and the other is not. The treatment effect for a specific individual i is given by

$$TE_i \equiv \hat{e}_{i2}(p_{i2}(1), c_{i2}(1)) - \hat{e}_{i2}(p_{i2}(0), c_{i2}(0)).$$

From Theorem 1, the linear-quadratic assumption implies that

$$TE_i \equiv \Gamma_c(c_{i2}(0) - c_{i2}(1)) + \Gamma_\ell(p_{i2}(1) - p_{i2}(0)).$$

Under the limited heterogeneity assumption, this can be written as

$$TE_i \equiv \Gamma_c(\mathcal{C}(0) - \mathcal{C}(1)) + \Gamma_\ell(p_{i2}(1) - p_{i2}(0)).$$

We will call $PE \equiv \Gamma_c(\mathcal{C}(1) - \mathcal{C}(0))$ the “*preference effect*.” Under the limited heterogeneity assumption, this effect is the same for all of the agents. In contrast, the *learning effect* $LE_i \equiv \Gamma_\ell(p_{i2}(1) - p_{i2}(0))$ varies across the agents. This is true even if they all have the same prior $p_{i1} = p_1$, because the posterior beliefs depend on effort, and with heterogeneous first-period costs c_{i1} , the first-period efforts \hat{e}_{i1} will also be heterogeneous.

One goal of the econometrician is to estimate the preference effect. As we will see, to do this it will be helpful to also estimate or at least bound the learning effect. To this end, let $\mathcal{X}_0, \mathcal{X}_1$ be two events corresponding to realizations of the assignment variable x_{i1} , with the assignment variable negative in \mathcal{X}_0 , so that $D_{i1} = 0$, and positive in \mathcal{X}_1 , so $D_{i1} = 1$. In what follows we will consider various definitions of these sets.

For events $\mathcal{X}_0, \mathcal{X}_1$ we define the *weighted average learning effect* by¹⁵

$$WAlE \equiv \Gamma_\ell (\mathbb{E} [p_{i2}(1) - p_{i1} \mid \mathcal{X}_1] - \mathbb{E} [p_{i2}(0) - p_{i1} \mid \mathcal{X}_0]).$$

This is the expected difference between the learning effect on the two sets of agents, where the conditioning includes the fact that agents with different priors and/or costs will have chosen different levels of first-period effort and so a) face different probability distributions over the assignment variable and b) update their beliefs for a given assignment in different ways. Because higher effort makes better outcomes more likely, and treatment leads agents to expend higher second-period effort, both of these effects go in the same direction, so $WAlE$ is always strictly positive. To bound its size, we use the *average prior strength*

$$\kappa_j \equiv \mathbb{E} [|1 - 2p_{i1}| \mid \mathcal{X}_j],$$

which again incorporates the selection effect of diverse priors and costs.

Combining the definition of the $WAlE$ with Theorem 2 yields

Theorem 3.

$$0 < WAlE < \frac{\Gamma_\ell}{\min\{F(-\bar{\gamma}), 1 - F(0)\}} (2 - \kappa_1 - \kappa_0).$$

Proof. Part 1 of Theorem 2 shows that $p_{i2}(1) - p_{i1} > 0$. Combined with Part 4 for $D_{i1} = 1$ this is

$$0 < p_{i2}(1) - p_{i1} \leq \frac{1}{\min\{F(-\bar{\gamma}), 1 - F(0)\}} (1 - \kappa_{i1}).$$

Part 1 of Theorem 2 also shows that $p_{i2}(0) - p_{i1} < 0$. Combined with Part 4 for $D_{i1} = 0$ this is

$$0 < -(p_{i2}(0) - p_{i1}) = |p_{i2}(0) - p_{i1}| \leq \frac{1}{\min\{F(-\bar{\gamma}), 1 - F(0)\}} (1 - \kappa_{i1}).$$

Hence

$$\begin{aligned} 0 < \mathbb{E} [p_{i2}(1) - p_{i1} \mid \mathcal{X}_1] &\leq \frac{1}{\min\{F(-\bar{\gamma}), 1 - F(0)\}} (1 - \kappa_1) \\ 0 < -\mathbb{E} [p_{i2}(0) - p_{i1} \mid \mathcal{X}_0] &\leq \frac{1}{\min\{F(-\bar{\gamma}), 1 - F(0)\}} (1 - \kappa_0). \end{aligned}$$

□

¹⁵To lighten notation, we do not write $WAlE(\mathcal{X}_0, \mathcal{X}_1)$.

This theorem shows that the $WAlE$ will be small when agents have high confidence in whether the state is $\bar{\gamma}$ or $\underline{\gamma}$ (so that κ_0 and κ_1 are near 1), or when the state does not much matter (so that $\bar{\gamma} - \underline{\gamma}$ is small). In either of these cases, the agents have a good sense of the effectiveness of effort, and so learning does not have much impact on their choice of effort in the second period. As we will see, difference-in-differences or regression discontinuity estimates of TE will then both provide good approximations of the PE , but the estimates can be far from the PE when the $WAlE$ is large.

2.1. Estimation

To estimate the treatment, preference, and learning effects, the econometrician uses observations of the second period outcome y_{i2} . The econometrician also observes either y_{i1} or x_{i1} . In the former case the econometrician uses difference-in-differences to estimate an average treatment effect; in the latter case they use a regression discontinuity analysis to estimate a local average treatment effect. If the econometrician observes both they may use either or both separately. Here we determine what is estimated by these two different techniques.

Regression Discontinuity

Suppose that the econometrician observes each agent's assignment variable x_{i1} , and wishes to compare the second-period behavior of agents whose assignment variable is just above or below the cutoff. To analyze this, let $\mathcal{X}_1^a = [0, a]$ and $\mathcal{X}_0^a = [-a, 0]$. With enough data the econometrician can consistently estimate $\mathbb{E}[y_2 | \mathcal{X}_1^a]$ and $\mathbb{E}[y_2 | \mathcal{X}_0^a]$. Since $y_{it} = e_{it} + \epsilon_{it}$ where the ϵ_{it} have zero mean and are independent of all other shocks and from Theorem 1 $\hat{e}_{it} = \Gamma_0 - \Gamma_c c_{it} + \Gamma_\ell p_{it}$, iterated expectations implies

$$\mathbb{E}[y_{i2} | \mathcal{X}_0^a] = \Gamma_0 - \Gamma_c \mathbb{E}[c_{i1} | \mathcal{X}_0^a] - \Gamma_c \mathcal{C}(0) + \Gamma_\ell \mathbb{E}[p_{i2}(0) | \mathcal{X}_0^a]$$

and

$$\mathbb{E}[y_{i2} | \mathcal{X}_1^a] = \Gamma_0 - \Gamma_c \mathbb{E}[c_{i1} | \mathcal{X}_1^a] - \Gamma_c \mathcal{C}(1) + \Gamma_\ell \mathbb{E}[p_{i2}(1) | \mathcal{X}_1^a].$$

It follows that

$$\mathbb{E}[y_{i2} | \mathcal{X}_1^a] - \mathbb{E}[y_{i2} | \mathcal{X}_0^a] =$$

$$PE - \Gamma_c (\mathbb{E}[c_{i1} | \mathcal{X}_1^a] - \mathbb{E}[c_{i1} | \mathcal{X}_0^a]) + \Gamma_\ell (\mathbb{E}[p_{i2}(1) | \mathcal{X}_1^a] - \mathbb{E}[p_{i2}(0) | \mathcal{X}_0^a]).$$

Now let

$$WAlE(a) \equiv \Gamma_\ell (\mathbb{E}[p_{i2}(1) - p_{i1} | \mathcal{X}_1^a] - \mathbb{E}[p_{i2}(0) - p_{i1} | \mathcal{X}_0^a]).$$

Because the c_{i1} and p_{i1} have a continuous joint density, and effort is a continuous function of cost and prior beliefs, these expectations are continuous in a , so in the limit as $a \rightarrow 0$ the two averages whose difference defines the $WAlE$ are taken with respect to the same weights. Thus $WAlE = \lim_{a \rightarrow 0} WAlE(a)$ is a consistent local estimate of the learning effect LE .

Theorem 4. *The econometrician can consistently estimate the local average treatment effect,*

$$LATE = PE + WAlE.$$

Thus our critique of the regression discontinuity approach is not that it gives an inconsistent estimate of the LATE, but rather that it might not be the real effect of interest.¹⁶

Difference-in-Differences

Suppose that the econometrician observes the outcomes y_{it} in both periods. As these are noisy signals of each agent's effort in each period, she can apply a difference-in-differences estimator. In this case

$$y_{i2} - y_{i1} = \hat{e}_{i2} - \hat{e}_{i1} + \epsilon_{2i} - \epsilon_{1i} = \Gamma_c(\mathcal{C}(D_{i1}) + \epsilon_i) + \Gamma_\ell(p_{i2}(D_{1i}) - p_{i1}) + \epsilon_{i2} - \epsilon_{i1}.$$

Hence the subsample mean $\bar{m}(D_1)$ is a consistent estimate of

$$\mathbb{E}[\Gamma_c \mathcal{C}(D_1) + \Gamma_\ell(p_{i2}(D_1) - p_{i1}) \mid \mathcal{X}_{D_1}]$$

where $\mathcal{X}_0 = \{x_{i1} < 0\}$ and $\mathcal{X}_1 = \{x_{i1} > 0\}$.

Theorem 5. *The difference in subsample means $\bar{m}(1) - \bar{m}(0)$ is a consistent estimator¹⁷ of the difference-in-differences effect,*

$$DIDE \equiv PE + WAlE,$$

where as above the $WAlE$ is with respect to $\mathcal{X}_0 = \{x_{i1} < 0\}$ and $\mathcal{X}_1 = \{x_{i1} > 0\}$.

¹⁶Deaton [6] makes a similar criticism of regression discontinuities for a different reason, that of heterogeneous treatment effects

¹⁷Unmeasured characteristics 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.

To understand better what the *WAlE* means in difference-in-differences estimation, consider first the case where all agents have the same prior $p_{i1} = p_1$ about the distribution of γ . In this case the difference between the groups \mathcal{X}_0 and \mathcal{X}_1 comes about solely because of the differential effort induced by different first-period costs, and on average agents in \mathcal{X}_0 will have lower effort levels than agents in \mathcal{X}_1 . From Theorem 2 we can use this to analyze the bias in the *DIDE*. Define for $j = 0, 1$

$$ALE_j \equiv \mathbb{E} [\Gamma_\ell (\mathbb{E} [p_{i2}(1) - p_{i1} \mid \mathcal{X}_j] - \mathbb{E} [p_{i2}(0) - p_{i1} \mid \mathcal{X}_j])]$$

where $\mathbb{E} [p_{i2}(1) - p_{i1} \mid \mathcal{X}_0]$ computes the conditional expectation of the counterfactual beliefs that individuals in \mathcal{X}_0 would have had had they observed success instead of failure and $\mathbb{E} [p_{i2}(0) - p_{i1} \mid \mathcal{X}_1]$ is the corresponding counterfactual for individuals who succeeded. Correspondingly, let $ATE_j \equiv PE + ALE_j$.

Theorem 6. *When all agents have the same prior, $ALE_1 \geq WAlE \geq ALE_0$ so $ATE_1 \geq DIDE \geq ATE_0$.*

In other words, the *DIDE* lies in between the average treatment effect for the two groups and may be viewed as a reasonable compromise between the larger and smaller effects. However, as with regression discontinuity, our main point is not about the accuracy of the estimated effect but rather a caution about its relevance.

Proof. Because the assignment variable in \mathcal{X}_1 stochastically dominates that in \mathcal{X}_0 , and the probability that $x_{i1} > 0$ is monotone increasing in first-period effort, it follows that the same is true for the effort levels. From Theorem 2 it then follows that the posterior $p_{i2}(1)$ in \mathcal{X}_1 stochastically dominates that in \mathcal{X}_0 and the posterior $p_{i2}(0)$ in \mathcal{X}_0 stochastically dominates that in \mathcal{X}_1 . Hence

$$\mathbb{E} [p_{i2}(1) - p_1 \mid \mathcal{X}_1] \geq \mathbb{E} [p_{i2}(1) - p_1 \mid \mathcal{X}_0]$$

and

$$\mathbb{E} [p_{i2}(0) - p_1 \mid \mathcal{X}_0] \geq \mathbb{E} [p_{i2}(0) - p_1 \mid \mathcal{X}_1].$$

Hence

$$\begin{aligned} ALE_0 &\equiv \mathbb{E} [p_{i2}(1) - p_1 \mid \mathcal{X}_0] - \mathbb{E} [p_{i2}(0) - p_1 \mid \mathcal{X}_0] \\ &\leq \mathbb{E} [p_{i2}(1) - p_1 \mid \mathcal{X}_1] - \mathbb{E} [p_{i2}(0) - p_1 \mid \mathcal{X}_0] = WAlE \end{aligned}$$

and

$$\begin{aligned} ALE_1 &\equiv \mathbb{E} [p_{i2}(1) - p_1 \mid \mathcal{X}_1] - \mathbb{E} [p_{i2}(0) - p_1 \mid \mathcal{X}_1] \\ &\geq \mathbb{E} [p_{i2}(1) - p_1 \mid \mathcal{X}_1] - \mathbb{E} [p_{i2}(0) - p_1 \mid \mathcal{X}_0] = WAlE, \end{aligned}$$

giving the stated result. □

Theorem 6 assumes that all agents have the same prior beliefs. When instead the priors are heterogeneous, the picture is more complicated. Here the agents in group \mathcal{X}_0 will be more pessimistic about the effect of effort than those in group \mathcal{X}_1 . However, optimism and pessimism do not map directly onto the strength of the learning effect, as it is the degree of confidence κ_{i1} that matters. For example, suppose that initial costs are homogeneous and priors heterogeneous, with pessimists who are fairly sure effort does not matter, and optimists who are fairly sure that it does. Then the successes will mostly be optimists and the failures pessimists, and each sort of agent will be more likely to observe the treatment that they expected, which moves their priors less than when they see the opposite treatment. Hence the *WAlE* will be quite small compared to the *ALE* for each group, so the *WAlE* will understate the true learning effect. Section 4.4 explains why the conclusion of the theorem may also fail when the agent has a multi-dimensional learning problem.

2.2. Identification and Bias

We have showed that both regression discontinuity (*LATE*) and difference-in-differences (*DIDE*) are consistent estimators for the sum of the preference effect *PE* and a form of the weighted average learning effect *WAlE*. In neither case are the two effects separately identified. Theorem 3 shows that the *WAlE* is always positive, so that the *LATE* and *DIDE* always overestimate the preference effect *PE*. In particular, if they are negative then the *PE* is negative as well.

It is important to recognize that the *PE* and *WAlE* have rather different implications for policy: the former does not depend upon the agents' prior beliefs but the latter does. Suppose, for example, that empirical analysis established that the treatment has a positive effect in the sense that $PE + WAlE > 0$, but that this was due to a strong learning effect, and that the *PE* is actually negative. If the treatment were then 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. By contrast negative estimated values of $PE + WAlE$ are robust to learning and so are more useful for policy purposes than positive values.

The procedures described above will yield approximately correct estimates of *PE* when the learning effect is small. We consider several scenarios under which this seems likely to be the case.

Success or Failure Uninformative

The learning effect will vanish for the regression discontinuity estimate if the agents observe x_{i1} , as they may if the underlying score variable is a test score or the

vote differential in an election. In this case, their success or failure gives agents no additional information about the efficacy of effort beyond the information contained in x_{i1} , so their posterior is $p_{i2}(x_{i1})$ independent of D_{i1} . As p_{i2} is continuous in x_{i1} , this implies that near the threshold where $x_{i1} = 0$ there is no discontinuity due to learning, and the regression discontinuity estimate will be equal to PE . Note that this logic need not apply to difference-in-differences estimation. While it is true that beliefs do not depend on D_{i1} , they do depend on x_{i1} . If x_{i1} is observed by the econometrician as well as by the agent, it can be included in the regression, which eliminates the learning effect. However, if x_{i1} is not included in the regression, then D_{i1} will proxy for x_{i1} , and there will still be a bias due to the learning effect.

Long Time Lag

The learning effect arises because with a small sample, idiosyncratic differences in information between different agents matter. As time goes on and samples grow larger, estimates will converge and the learning effect will diminish. Hence if a long period elapses between the signal of success or failure and the second period effort choice agents may have acquired additional information about the effect of effort and the learning effect will be attenuated.

There are two caveats. First, it should be noted that if effort has a large effect ($\gamma = \bar{\gamma}$), then the losers will reach the wrong conclusion and reduce their effort. This means that they will get less feedback about whether effort matters, so it may take a long time before they learn the truth. Second, learning may have a persistent indirect effect, for example if effort is an investment.

3. Applications

3.1. 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_{it} corresponds to resisting the insurgents, x_{i1} is the score the Air Force assigned to village i , and $D_{it} = 0$ corresponds to village i being bombed. One theory is that bombing angers villagers, making them less inclined to resist, so $\mathcal{C}(0) > \mathcal{C}(1)$. A competing theory is that bombing destroys infrastructure and makes the village less attractive to insurgents so $\mathcal{C}(0) < \mathcal{C}(1)$.

Dell and Querubin [7] 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 initial scores that researchers observed (with some difficulty) but villagers

did not. Villages were then put into a small number of categories based on their scores and bombed (or not) accordingly.¹⁸ The researchers also observed several different outcome measures y_{i2} of subsequent effort, including an *ex post* set of scores. Dell and Querubin [7] compares the outcomes of villages with initial scores just below the threshold with the outcomes of villages that were just above it to estimate the treatment effect of bombing by regression discontinuity analysis, and find that “[b]ombing increased the military and political activities of the communist insurgency, weakened local government administration and lowered non-communist civic engagement.”¹⁹

Our analysis suggests that the negative treatment effect of bombing on compliance that Dell and Querubin [7] identified may have been due to learning, and that the preference effect may well have been 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 leads to a negative estimated effect if the preference effect is zero, and can lead to negative estimates despite a positive preference effect if the villagers were sufficiently uncertain about how much effort mattered.

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 the U.S. dropped leaflets to advertise that effort matters (Dell and Querubin [7]), which might have led villagers to believe that effort matters, but the leaflets did not answer the question of how much effort matters. 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 [7] whether the induced change was the result of a change in preferences or a more ephemeral effect of learning.

To assess the preference 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 [20] 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

¹⁸Our simple model has only a single threshold; the next section briefly discusses the extension to multiple thresholds.

¹⁹They 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.

drinking heavily. Of course the mere fact that targeting was indiscriminate does not indicate that villagers knew this, but Lyall [20] 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 the villagers made formal complaints about being targeted for no reason. Thus in the Lyall [20] study there is reason to believe that learning effects were small. Here the outcome measure of effort y_{it} was insurgent activity with y_{i1} before the shelling and y_{i2} after the shelling.²⁰ Lyall [20] estimates the effect of shelling using difference-in-differences, and finds that insurgent activity decreased after shelling, the opposite conclusion from Dell and Querubin [7].

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 effort matters. 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.

3.2. 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 treatment $D_{i1} = 1$ corresponds to winning the election, and $D_{i1} = 0$ to losing it. Effort corresponds to campaign effort. Here the treatment corresponds to $D_{i1} = 1$ 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, that is $\mathcal{C}(1) < \mathcal{C}(0)$, so it should increase effort and the 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 find it costly to campaign in her local district.

²⁰Strictly speaking, if the shelling was indiscriminate difference-in-differences is not needed, as the treated and untreated villages should be the same. Lyall [20] 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.

This issue was studied by Lee [16]. The outcome measure y_{i2} was the vote share in the subsequent election and the data was analyzed using a regression discontinuity analysis. Here the latent variable x_{i1} is the vote differential and naturally those who just lose an election are not so different from those who just win. Lee [16] finds a substantial discontinuity indicating that just barely winning greatly increases the chances of future success.

As in Dell and Querubin [7], effort provision here is endogenous and depends upon beliefs about the efficacy of effort in turning out votes. Moreover, this goes in the “wrong” direction for Lee: The learning effect means that winners would provide more effort and losers less, so a positive effect could in principle be due to learning and not a true incumbency effect through the cost function. However, there is reason to believe that the learning effect is small here. First, political parties have a lot of data about past elections, so they probably have a pretty good idea about the efficacy of effort. Second, even if this is not the case, the latent variable x_{i1} is directly observed, so the actual success or failure conveys no additional information, and we predict the learning effect will not be present.²¹

3.3. Education

In the educational setting there is a great deal of interest in whether attending a prestigious school is beneficial, for example, if it increases future earnings. Here we suppose that the agents are applicants, and that the treatment $D_{i1} = 1$ corresponds to being admitted to a prestigious school, so that the x_{i1} are the test scores used to determine admission. Here the initial value c_{i1} represents the applicant’s intrinsic ability, with higher ability corresponding to a lower cost of effort. The preference effect $\mathcal{C}(1) - \mathcal{C}(0)$ represents the value added from attending the prestige school.

Hoekstra [15] studied this setting with the outcome measure y_{i2} of effort being earnings roughly ten years after the admissions decision, and analyzed the data using a regression discontinuity analysis. The paper says that

“the admission rule was never published or revealed by the university and, in fact, was changed (albeit moderately) from year to year. Consequently, it is unlikely that the applicant would know, prior to applying, whether she was just above the cutoff or just below it.”

²¹Note 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 does not seem relevant in the elections Lee [16] studies, as the candidates presumably knew the rules. Note also that candidates may well have information also about other elections, but as we do not expect a learning effect in this setting, it is not relevant.

Moreover, individual students may often not know very much about the effort levels of other students. As a consequence, this study is subject to a possible learning effect, because being admitted or rejected conveyed information about the effectiveness of effort. However, the paper finds a strong effect after nearly a decade: Those who just made it over the line had about 20% higher earnings than those who just missed, and it is implausible that a direct learning effect this strong would persist over such a long period of time. An indirect effect might have greater persistence, because effort in schooling represents an investment, and the timing of that investment is important. Hence the losers might eventually learn the truth, but by the time they do they will have acquired less human capital than the winners, and it may be quite costly to make up the difference.

Note also that if it was revealed to applicants how close they were to the cutoff, then admission would no longer signal that “effort matters.” Hence if we could show that the estimated impact on earnings is robust to the informational condition, we could reach a stronger conclusion. The strength of the learning effect is also important for assessing changes in admission policies.

To see how the learning effect matters in the setting of our baseline model, consider an admissions policy that lowers the threshold for success for a disadvantaged group and that applicants know by how much the threshold is changed. To understand the decomposition of the treatment effect into a preference and learning effect we consider the impact of the change on the disadvantaged group in two opposite cases: Either there is only a preference effect or there is only a learning effect.²²

If there is only a preference effect, then more disadvantaged applicants receive the benefit of the preference effect $\mathcal{C}(1) - \mathcal{C}(0)$: these applicants have lower effort cost in the second period, provide more effort and achieve greater earnings y_{i2} . If there is only a learning effect, the impact on disadvantaged applicants depends on whether their status is changed by the change in threshold. Recall that the decrease in threshold is equivalent to lowering F_0 , decreasing the probability of failure. Under the linear-quadratic assumption, this does not change the marginal incentives for first period effort, so by Theorem 1 first period effort provision does not change. Hence for applicants who are admitted under the new policy but would not have been under the old there is a benefit from the learning effect: They achieve success and their posterior beliefs increase, rather than failing and having decreased posterior beliefs. Thus these admittees have a greater second period effort and higher earnings

²²Here we assume the preference effect benefits of admission are not changed by the change in policy, and do not consider how it might impact members of other groups. Applicants whose thresholds are raised would face the opposite effects.

y_{i2} than they would have had in the absence of the program. The effect of lowering the threshold is different for disadvantaged applicants who would not have been admitted under either threshold, and also for those who would have been admitted under both thresholds. From Theorem 2 part (3), their posteriors are lower under the new threshold than under the old one. This means that both groups provide lower effort and have lower second period earnings y_{i2} , so both of these groups are harmed. Thus to evaluate the overall effect of the change on the disadvantaged applicants, we must weight the beneficial effect of learning on the disadvantaged applicants who are admitted only because of the new threshold against the negative impact on the other disadvantaged applicants.

Of course this analysis presumes that the applicants are only uncertain about the effectiveness of effort and not about the threshold itself, so that they know the probability of being admitted when they take no effort at all. Section 4.4 discusses the more complex case where both the effectiveness of effort and the cutoff are both unknown.

4. Discussion

4.1. Implications for Policy

What is the relevance of the estimated treatment effects for policy decisions about which treatments to implement? Note first that the benefit of the treatment is not simply that the treated group behaves differently from the non-treated group. For example, admission to an exclusive college as a reward may induce high school students to study more if they extrapolate their college application experience to the rewards to effort in college. In other words, part of the benefit of treatment is that it provides incentives for effort both before and after treatment. This implies policy evaluations may have to weigh increased effort due to incentives against decreased effort *ex-post* due to failure. 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.

Second, it clearly matters for policy whether the treatment effect is due to a change in preferences or to learning. Suppose that the treatment reduces *ex post* effort. It seems then as if there is a tradeoff between the reduced effort due to failure and the increased effort due to incentives. If the effect of the treatment is solely due to a change in preferences this is correct. Suppose, however, that most or all of the reduction in *ex post* effort is due to learning. Here, although it is true that

agents who fail lower their effort level, they still provide a higher level of effort than they would have in the absence of incentives. Hence if the treatment effect is due to learning there is in fact no trade off between incentives and the treatment effect, and the estimated treatment effect is not relevant to the evaluation of the policy.

4.2. Time Lags

In some of the settings we would like to consider, the measurement may be taken long after second period effort is made, or both effort and measurement may take place long after the treatment. In other settings, there may be more than two periods, for example a third election after the second, a third round of bombing after the second, and so forth. We will briefly indicate some of the considerations that impact the learning effect.

As we noted in Section 2.2, if the outcome occurs long after the treatment we would not expect the direct learning effect to be important: as time goes on additional data beyond the treatment will be acquired and the learning effect will diminish. For example, in the Dell et al. [8] study the effect of being conquered by the Khmer 150 years ago on current economic prosperity in South Vietnam is unlikely to have a substantial learning component. To loosely relate this to our model, suppose that the effort used to deter takeovers by outsiders is detrimental to economic development. Then villagers who resisted invasion would infer that effort is effective and so work harder at deterrence, which could lower their economic well-being. It seems unlikely that this sort of learning would have a persistent direct effect, though it is conceivable that it could have a persistent indirect effect through its effect on culture and institutions.

The situation with respect to preferences is more complicated and context dependent. We might expect that the preference effect also declines over time; for example the effect on preferences of a bombing that took place several days ago might well be greater than a bombing that took place several decades ago. Second, the impact of additional “treatments” beyond the first one might attenuate as well. In the case of incumbency, winning additional elections might have relatively little effect. On the other hand, while bombing or shelling a second time might have less impact than the first time, it could still be substantial.

Finally, we have assumed that the agents’ second period problem is similar to their problem in the first period, but over time the incentive structure might well change. In the case of state violence, for example, the incentives are designed to influence a conflict that will eventually end, and we may be primarily interested in the long-term effects of the treatment after the conflict has ended, when the effort incentives as measured by γ will no longer be present. This should lead to a reduction of effort

for both the treated and control groups, and also means that inferences about γ based on first period data are irrelevant: Everyone knows that there will be no more bombing after the war ends, regardless of effort. Nevertheless the preference effect may persist, and we may wonder whether it does. This has implications for data gathering. For example in Dell and Querubin [7] it might not have been practical to get post-war data on villages, but it might be possible to focus on data about effort gathered near the end of the war, when it was obvious that the Americans were going to lose and the bombing would stop.

4.3. 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 x_{i1} 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 at all.

The presence of multiple thresholds potentially complicates analysis as it can lead to non-convexities in the agent’s objective function. Here we assume that the 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 students 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.

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.²³

²³In Dell and Querubin [7] 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.

4.4. 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 extent to which effort lowers the probability of failure. In our baseline model, they are however certain about the probability of success when no effort is made. More generally, they might also be uncertain about the probability of success when they take no effort, perhaps because they are uncertain of the threshold used to determine success. To model this, suppose that the effect of effort is given by $\chi + \gamma e_{it} + \omega_{it}$ where χ may take on one of two values $\underline{\chi}, \bar{\chi}$ with $\bar{\chi} > \underline{\chi}$. Here $\bar{\chi}$ represents the possibility that success is relatively likely, for example, that a test has a low passing score. The agent is uncertain about both χ and γ , so their beliefs have the form $p_{it}(\chi, \gamma)$.

We strengthen the linear quadratic assumption so that F is linear in ω over the range $[-1 - \bar{\chi}, -\underline{\chi}]$. In this case χ does not have any direct effect on effort, so Theorem 1 remains correct as stated with respect to the induced marginal distribution \tilde{p}_{i1} over γ . Suppose, moreover, that we make the normalization that the prior expected value of χ is zero, that is $\tilde{p}_{i1}(\bar{\chi})\bar{\chi} + \tilde{p}_{i1}(\underline{\chi})\underline{\chi} = 0$. If the prior is independent between χ and γ , so that $p_{i1}(\chi, \gamma) = \tilde{p}_{i1}(\chi)\tilde{p}_{i1}(\gamma)$, then as we show in the online Appendix Theorem 2 is also correct as stated.

With independence in the priors, then, nothing changes. However, there may be negative correlation between the slope and intercept: If the probability of failure under no effort is very low, then increasing effort cannot decrease it much further, so χ and γ might be negatively correlated. To see how negative correlation matters, consider the simple case where there is perfect negative correlation, so the only two possibilities are $\bar{\chi}, \underline{\gamma}$ and $\underline{\chi}, \bar{\gamma}$. Consider the intersection $\bar{\chi} + \underline{\gamma}e^* = \underline{\chi} + \bar{\gamma}e^*$. For higher levels of effort $e_{it} > e^*$ we have $\underline{\chi} + \bar{\gamma}e_{it} > \bar{\chi} + \underline{\gamma}e^*$ meaning that success is more likely with $\bar{\gamma}$ than with $\underline{\gamma}$. As before, success is a signal that effort matters, so it increases effort, while failure lowers effort. For lower levels of effort, however, we have $\underline{\chi} + \bar{\gamma}e_{it} < \bar{\chi} + \underline{\gamma}e^*$, so for low effort levels success is a signal that effort does not matter. As an intuitive 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 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.

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 [7] and Lee [16] examples discussed above, especially in Dell and Querubin [7] where most of the data is for high-effort villages. However, this may not be true in other contexts.

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 preference effect.

4.5. *Non-Myopic Agents*

If agents are not myopic they will consider the effect of their first-period effort on their expected second-period value. Let $\hat{v}(p_{it}, c_{it}) = \max_{e_{it}} v(e_{it} | p_{it}, c_{it})$ represent the myopic utility from the optimal effort decision. Letting $p_{i2}(D_{i1}|e_{i1})$ represent the posterior as a function of the outcome and the first period effort the expected second period value is given by

$$V_{i2}(e_{i1}) = p_{i2}(e_{i1})\hat{v}(p_{i2}(1|e_{i1}), c_{i2}(1)) + (1 - p_{i2}(e_{i1}))\hat{v}(p_{i2}(0|e_{i1}), c_{i2}(0)).$$

This complicates the first period problem because $V_{i2}(e_{i1})$ is not quadratic in e_{i1} and indeed may even fail to be concave. Nevertheless V_{i2} is continuous, so there is an optimal first period choice of effort, but this choice is different when the agent is patient rather than myopic. We now explain how this impacts our results.

Theorem 2 does not depend upon how first period effort is chosen. Our results on regression discontinuity are not based on first period observations at all, but only on second period comparisons, so they are valid regardless of how first period effort is chosen. The situation for difference-in-differences is more complicated. We already observed that the waters are muddied if beliefs are heterogeneous. When the agent is patient, heterogeneity in cost also muddies the waters. The reason we get clean results with homogeneous beliefs and heterogeneity only in costs is that first period effort is linear in first period cost, so cancels out of the calculation. Once first period effort is non-linear and perhaps even non-monotone in first period cost, difference-in-differences faces an additional confounding factor in that the first and second period cost difference effects no longer cancel out.

Finally, we can say something about the effect on first period effort of relaxing the myopia assumption. If we use the envelope theorem to compute the derivative of second period utility with respect to effort we find

$$V'_{i2}(e_{i1}) = [(\hat{v}(p_{i2}(1|e_{i1}), c_{i2}(1)) - \hat{v}(p_{i2}(0|e_{i1}), c_{i2}(0))) p'(1|e_1)] + \\ [f(\bar{\gamma} - \underline{\gamma}) [e_{i2}(1) - e_{i2}(0)] p(1|e_1) + f(\bar{\gamma} - \underline{\gamma}) e_{i2}(0)].$$

Here there are two effects. The first effect is due to the impact of first-period effort increasing the probability of success. If success lowers second period costs it leads to a positive utility differential and this effect is positive, otherwise it might not be. The second effect is the informational effect: Increasing effort makes the signal more informative, which improves second period decision making. Since the optimal first-period effort for a myopic agent is given by the unique solution of the first-order condition, strictly increasing the derivative by adding the positive expression $\delta V'_{i2}(e_{i1})$ (where δ is the discount factor) necessarily shifts the optimum to the right. Thus, when success lowers second period costs a non-myopic agent exerts greater effort, and from Theorem 2 the learning effect is stronger.

4.6. Imperfectly Correlated Effectiveness of Effort

We have assumed that the effectiveness γ of effort is exactly the same in each period. More generally, there might be two different parameters γ_1, γ_2 in the two periods. In the extreme case where these are independent, learning about γ_1 does not change beliefs about γ_2 , and the learning effect vanishes. However, if the binary random variables γ_1, γ_2 are positively correlated, the learning effect will be positive as well, though smaller than in our case of perfect correlation. For example, if first period effort describes effort in high school, and second period effort describes effort in college, then agents might believe that there is imperfect positive correlation in the γ 's, which would reduce the gap between the estimated treatment effect and the preference effect.

4.7. Learning From Others

In our learning model, agents learn only from their own experience. In practice, agents may learn from the experiences of others as well: villagers may learn about which other villages have been bombed, and students may learn about the efforts and outcomes of friends and relations. If this learning takes place before the initial effort decision then the agents should have fairly tight priors. It follows from Theorem 3 that the *WALE* will be small, so both for regression discontinuity and difference-in-difference the estimated treatment effect will be close to the preference effect. On

the other hand, if agents learn from the experience of others after the initial effort decision has already been taken, we expect the learning effect to be larger, because a larger number of conditionally independent signals should lead to more variation in posterior beliefs.

5. 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, better instruments, and develop improved techniques. 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. Specifically, when selection into the treatment depends on the agents' effort, being treated may provide the agents with information that influences their behavior. As we have shown, this can lead to quite different policy recommendations than when the 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 does not add information to the vote differential. To make our points most clearly, we have made strong functional form assumptions, but we expect the points that we have articulated here to apply much more generally.

References

- [1] Alberto Alesina, Edward Glaeser, and Bruce Sacerdote. Why the US Doesn't Have a European-Style Social Welfare System. *NBER Working Papers*, 8524, 2001. ISSN 1098-6596. doi: 10.1017/CBO9781107415324.004.
- [2] Joshua David. Angrist and Jörn-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=ztXL21X
- [3] Sylvain Chassang, Padró I Miquel, Erik Snowberg, et al. Selective trials: A principal-agent approach to randomized controlled experiments. *American Economic Review*, 102(4):1279–1309, 2012.
- [4] Sylvain Chassang, Erik Snowberg, Ben Seymour, and Cayley Bowles. Accounting for behavior in treatment effects: New applications for blind trials. *PLoS One*, 10(6):e0127227, 2015.
- [5] Gilles Chemla and Christopher A Hennessy. Controls, belief updating, and bias in medical rcts. *Journal of Economic Theory*, 184:104929, 2019.
- [6] Angus Deaton. Instruments, randomization, and learning about development. *Journal of Economic Literature*, 48(2):424–55, 2010.
- [7] 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.
- [8] Melissa Dell, Nathan Lane, and Pablo Querubin. The historical state, local collective action, and economic development in vietnam. *Econometrica*, 86(6):2083–2121, 2018. doi: 10.3982/ECTA15122. URL <https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA15122>.
- [9] Rohan Dutta, David K Levine, and Salvatore Modica. Interventions when Social Norms are Endogenous: A Critique. 2018.
- [10] Drew Fudenberg and David K Levine. Self-Confirming Equilibrium. *Econometrica*, 61(3):523–545, 1993.

- [11] Drew Fudenberg and David K. Levine. Self-confirming Equilibrium and the Lucas Critique. *Journal of Economic Theory*, 144(6):2354–2371, 2009. ISSN 00220531. doi: 10.1016/j.jet.2008.07.007. URL <http://dx.doi.org/10.1016/j.jet.2008.07.007>.
- [12] Andrew Gelman and Gary King. Estimating incumbency advantage without bias. *American Journal of Political Science*, pages 1142–1164, 1990.
- [13] James J Heckman. Causal parameters and policy analysis in economics: A twentieth century retrospective. *The Quarterly Journal of Economics*, 115(1): 45–97, 2000.
- [14] James J Heckman and Jeffrey A Smith. Assessing the case for social experiments. *Journal of Economic Perspectives*, 9(2):85–110, 1995.
- [15] Mark Hoekstra. The effect of attending the flagship state university on earnings: A discontinuity-based approach. *The Review of Economics and Statistics*, 91(4):717–724, 2009. doi: 10.1162/rest.91.4.717. URL <https://doi.org/10.1162/rest.91.4.717>.
- [16] 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.
- [17] David S. Lee and Thomas Lemieux. Regression Discontinuity Design in Economics. *The Journal of Economic Literature*, 48:281–355, 2008. ISSN 1536867X. doi: 10.1257/jel.48.2.281.
- [18] Edwin Leuven, Lars J. Kirkeboen, and Magne Mogstad. Field of Study, Earnings, and Self-Selection. *The Quarterly Journal of Economics*, 131(3): 1057–1111, 05 2016. ISSN 0033-5533. doi: 10.1093/qje/qjw019. URL <https://dx.doi.org/10.1093/qje/qjw019>.
- [19] Robert Lucas. Econometric Policy Evaluation: A Critique. In Karl Brunner and Alan Meltzer, editors, *The Phillips Curve and Labor Markets*, pages 19–46. Carnegie-Rochester Conference Series on Public Policy, North-Holland, 1 edition, 1976.
- [20] Jason Lyall. Incite Insurgent Attacks? Evidence from Chechnya. *Journal of Conflict Resolution*, 53(2):331–362, 2009. doi: 10.1177/0022002708330881.

- [21] Tomas Philipson and Jeffrey DeSimone. Experiments and subject sampling. *Biometrika*, 84(3):619–630, 1997.
- [22] Mark R Rosenzweig and Kenneth I Wolpin. Natural “natural experiments” in economics. *Journal of Economic Literature*, 38(4):827–874, 2000.
- [23] Thomas Sargent, Noah Williams, and Tao Zha. Shocks and Government Beliefs : The Rise and Fall of American Inflation. *American Economic Review*, (1977): 1193–1224, 2006.
- [24] Christopher A Sims. Macroeconomics and methodology. *Journal of Economic Perspectives*, 10(1):105–120, 1996.
- [25] Donald L. Thistlethwaite and Donald T. Campbell. Regression-discontinuity Analysis: An Alternative to the ex post facto Experiment. *Journal of Educational Psychology*, 51(6):309–317, 1960. ISSN 0022-0663. doi: 10.1037/h0044319. URL <http://content.apa.org/journals/edu/51/6/309>.
- [26] John H. Tyler, John B. Willett, and Richard J. Murnane. Estimating the Labor Market Signaling Value of the GED. *The Quarterly Journal of Economics*, 115 (2):431–468, 05 2000. ISSN 0033-5533. doi: 10.1162/003355300554818. URL <https://dx.doi.org/10.1162/003355300554818>.
- [27] Halbert White. A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica: Journal of the Econometric Society*, pages 817–838, 1980.