The Paradox of Policy-Relevant RCTs and Natural Experiments*

Gilles Chemla
Imperial College, DRM/CNRS, and CEPR.

Christopher A.Hennessy
LBS, CEPR, and ECGI

April 2016

Abstract

According to conventional wisdom, RCTs and natural experiments represent especially credible bases for econometric inference, facilitating evidence-based policymaking. We assess credibility in dynamic settings, examining robustness of evidence derived from an exogenous first-stage randomization applied to measure zero subjects. If government is able (unable) to alter policy in response, experimental evidence is contaminated (uncontaminated) by *ex post policy endogeneity*: Measured responses depend upon the government objective function into which the evidence will be fed. Similarly, if government perceives experimental evidence as credible (non-credible), the very act of observation changes (does not change) agent behavior. Thus, paradoxically, the experimental evidence is contaminated if and only if government is willing and able to use it. Ex post endogeneity causes measured responses to hinge upon (unknown) parameters of the governmental objective function, as well as *a priori* beliefs regarding the causal parameters to be estimated. Moreover, heterogeneous causal effect parameters induce endogenous belief heterogeneity. This link between beliefs and causal effect parameters makes it difficult, and potentially impossible, to isolate the latter using experimental evidence. Finally, we show treatment-control differences in RCTs are contaminated unless stock variable accumulation cost functions satisfy strong functional form assumptions: zero fixed costs, equality of buy and sell prices, and quadratic adjustment costs.

*We thank seminar participants at Stanford, U.C. Berkeley, CMU, MIT, Boston College, Copenhagen, LBS, Imperial College, Washington-Seattle, Miami, and Nova-Lisbon. We also thank Antoinette Schoar, Manuel Adelino, and Taylor Begley for early suggestions. Hennessy acknowledges funding from the ERC.

1 Introduction

Randomized controlled trials (RCTs) are viewed as the gold standard for empirical evidence in medical science. Increasingly, this view is being carried over to the social sciences, with evidence from RCTs and natural policy experiments (NPEs) viewed by many as especially credible. For example, Greenstone (2009) writes, "The gold standard for estimating the causal impact of a regulation is the randomized trial." In recent years, econometricians have shown considerable ingenuity in attempts to exploit exogenous variation arising from NPEs, or in constructing their own RCTs. Angrist and Pischke (2010) herald these methods as a "credibility revolution," arguing empirical evidence delivered via quasi-randomization represents a credible stand-alone product. Their textbook (2009) states, "A principle that guides our discussion is that most of the estimators in common use have a simple interpretation that is not heavily model dependent."

Ingenuity notwithstanding, the credibility revolution has faced criticism for allegedly placing means before ends, elevating clever identification over economic relevance. In response, empirical researchers have redoubled their efforts to prove policy relevance. In fact, the perceived credibility of evidence derived from policy randomization has recently led to a greater willingness to make policy decisions based directly upon them. For example, Dhaliwal and Tulloch (2015) note the existence of an "increasing trend towards considering rigorous evidence while making policy decisions." Consistent with this view, the mission of J-PAL is to promote "evidence-based policymaking." In a similar vein, Greenstone (2009) calls for government agencies "to move toward a culture of persistent regulatory experimentation" in which randomized regulations are sunsetted so that impact analysis can inform the next regulatory decision. Duflo (2004) argues, "Creating a culture in which rigorous randomized evaluations are promoted, encouraged, and financed has the potential to revolutionize social policy during the 21st century, just as randomized trials revolutionized medicine during the 20th."

In this paper, we describe a potential internal inconsistency at the core of the natural experiment research program as it relates to inference in dynamic settings. A simple example illustrates. We begin by noting that the two primary, and often exclusive, objectives of contemporary empirical work are to convince the audience of "clean identification," through as-good-as-random assignment, and policy relevance. Suppose then that after exhaustive debate, which is the current norm in empirical due diligence, an empiricist is ultimately able to convince her audience that her identification strategy is "clean" in the sense that she can offer a compelling demonstration based upon her granular knowledge of institutional details that it was Nature herself that forced an exogenous change in government policy and/or that Nature herself randomly assigned agents to treatment and control groups. The empiricist is next challenged on policy-relevance. Suppose that here too, she

can rise to the challenge, demonstrating direct policy-relevance, e.g. the experimental outcomes are actually being utilized by the government in considering the scale-up of its pilot program. At this stage our empiricist is allowed to declare victory, and lauded for her careful and important study.

What has gone unnoticed in this parable seminar is that there is a potential contradiction between our empiricist's claim of clean identification, on one hand, and her demonstration of direct policy relevance on the other. After all, in establishing policy-relevance, our empiricist has also demonstrated that the probability distribution of the policy variable is being altered by the experimental evidence. But if agents are making forward-looking decisions, e.g. accumulating some stock variable or building a reputation, they will have rationally changed their behavior during the experiment in light of the anticipated influence of econometric analysis. In other words, the anticipation of evidenced-based policymaking post-experiment will change what the econometrician measures during the experiment. What implications does this have for causal inference?

To examine this question, we consider the following setting. Agents (firms) make optimal (investment) decisions at each instant in a (tractable) continuous-time infinite-horizon economy. They do so in light of current and expected future caps on pollution. Econometricians (and the government) would like to infer the marginal benefit each firm derives from increasing the pollution cap, since this determines the key causal effect parameter in the economy: the effect of the cap on long-term investment in polluting capital stocks. Marginal pollution benefits are private information to the firm. Each industry-specific pollution benefit represents an i.i.d. draw from a known distribution.

Fortunately for econometricians, ideal randomized evidence will arrive to shed light on the causal inference problem. The evidence takes one of two forms. In a *Natural Policy Experiment* (NPE), all firms will be subjected to a common exogenous shock to the pollution cap during a randomly-timed Experiment Stage. In a *Randomized Controlled Trial* (RCT), a fraction of firms will face regulation during the Experiment Stage and the remaining firms will not. In the NPE setting, econometricians attempt to infer industry-level pollution benefits based upon each industry's reaction to the regulation shock. In the RCT setting, inference will instead be predicated upon the difference between treatment and control group investment.

We explore whether and how experimental evidence is altered according to whether and how it will be used by the government. To this end, we consider three post-experiment scenarios. In the first scenario, the government is powerless to change the path of the policy variable post-experiment. In the second scenario, the government is able to change the path of the policy variable post-experiment, and will do so using the information generated by the experiment. For example, a pro-investment government might deregulate only if average industry-level pollution benefits are sufficiently high, with high benefits implying larger investment increases (causal effects) in response to deregulation. In the final scenario, the government has the power to change the policy variable

post-experiment, but will do so relying upon prior information, viewing experimental evidence as non-credible.

In this setting, feedback from experimental evidence to the probability distribution of the post-experiment policy variable contaminates the once-clean evidence. In particular, policy-relevant evidence from our seemingly ideal first-stage randomization is contaminated by what we term *ex post endogeneity*. And, as we show, this is true even if, as in our economy, agents are measure zero and have no strategic incentive to distort behavior under observation. But notice, the contamination of evidence that results from alteration of the future policy variable distribution necessarily vanishes if the government is unable to alter policy in the future. We thus have the following paradoxical situation: The experimental evidence is uncontaminated only if the government is unable to use it. Similarly, the contamination vanishes if the government does not view the experimental evidence as credible and so ignores it. We thus have another paradox: The experimental evidence is uncontaminated only if the government does not deem it to be credible.

Moving beyond these paradoxes, we show ex post endogeneity creates five novel challenges to inference. First, rather than being stand-alone objects, policy-relevant experimental evidence must be interpreted in light of the deep structural parameters of the governmental objective function into which the evidence will be fed. That is, the fact that one has observed a natural experiment in which the first-stage randomized policy does not depend on government objective function parameters does not eliminate the need to make assumptions about, and interpret evidence in light of, these parameters. After all, these parameters dictate the future distribution of the policy variable, and thus influence measured responses today. In other words, the endogeneity problem has been pushed back in time.

Second, we show that experimental evidence must also be interpreted in light of prior beliefs regarding the causal parameters to be estimated, even if, as in our setting, the true parameter values (industry-specific pollution benefits) will actually be inferred based on the econometric estimation, so that the prior beliefs are abandoned in policy-setting. That is, correct econometric inference is predicated upon *a priori* information. Intuitively, since agents know the government will base its post-experiment policy decisions on parameter estimates, prior beliefs regarding the probability distribution of these parameters map directly to prior beliefs regarding the distribution of the policy variable post-experiment, shaping the behavior of forward-looking agents during the experiment. An incorrect stipulation of prior beliefs regarding the distribution of causal parameters will lead to incorrect inference regarding the true value of these parameters.

Third, ex post endogeneity gives rise to John Henry and Hawthorne Effects ("observer effects"): Both control and treatment groups change their behavior under observation. And this is true despite our agents being rational, anonymous, and measure zero. Intuitively, variation in the post-

experiment policy variable distribution, resulting from econometricians' observation, changes behavior during the experiment.

Fourth, these observer effects are not equalized across treatment and control groups unless the underlying accumulation technology satisfies the type of "strong functional form assumptions" that randomization advocates have tended to criticize in structural work (e.g. Hayashi (1982)): zero fixed costs, equality of buy and sell prices, and quadratic adjustment costs. Absent these functional form assumptions, treatment-control differences are contaminated by observer effects.

Finally, as we show, under rational expectations, endowed heterogeneous causal effect parameters of agents generate *endogenous belief heterogeneity*, which then amplify or attenuate treatment response heterogeneity. The failure to make a correct accounting for the response amplification/attenuation resulting from heterogeneous beliefs would then lead to faulty inference regarding the magnitude of causal effect parameters. Worse still, in some cases this heterogeneous beliefs channel can make it impossible to extract causal parameters from RCT and NPE evidence. Formally, this is due to the fact that the heterogeneous beliefs channel can cause experimental moments, e.g. the treatment-control difference, to be non-monotone in causal effect parameters. For example, one can observe the same experimental moment under both high and low values of the true causal effect parameter, so the economic content of the experimental moment is no longer clear.

The intuition for the heterogeneous beliefs channel is as follows. Even under fully rational expectations with measure zero agents, those agents who benefit a great deal from an experimental government program, via their endowed causal effect parameters, would attach a higher probability to the government continuing the program based on the econometric evidence being fed into subsequent governmental cost-benefit analysis. This would amplify their response to the experimental program beyond what one would expect based on endowed causal effect parameters. Conversely, agents who are greatly harmed by, say, an experimental governmental regulation would attach a lower probability to the government continuing the regulation based on econometric evidence being fed into subsequent governmental cost-benefit analysis. But this would attenuate their response to the experimental regulation. Although we do not explore the question, such effects will be even more powerful for "behavioral agents" who place excessive weight on their own subjective experience.

Essential to the arguments above is that the agents exposed to the first-stage randomization make decisions during the experiment (as opposed to crops, cells, or particles), and that those decisions are functions of the future probability distribution of the policy variable the randomized experiment will influence. Although the paper considers the canonical example of cross-sectional real investment to fix ideas, our arguments apply to a range of decision problems involving the accumulation of stock variables: savings, debt, health, human capital, offspring, financial structure, employees, and reputation in principal-agent settings. Further, the arguments apply to many superficially "one-off"

decisions once one accounts for life-cycle considerations, e.g. labor supply. However, this is not to say that the problems flagged apply to all RCTs and NPEs. Essential for our arguments is that the agents being measured face the endogenous government policy responses to the newly-created evidence. Section 2 provides examples. However, we do not argue the problems flagged are universal. Rather, we simply argue that the problems highlighted will become increasingly relevant the tighter the estimation-policy setting nexus, a nexus that many randomization advocates promote.

The issues we raise are related to, but distinct from, the econometric critique made by Lucas (1976). Writing for The New Palgrave Dictionary of Economics, Lars Ljungqvist (2008) offers the following definition of the Lucas Critique:

> It criticizes using estimated statistical relationships from past data to forecast effects of adopting a new policy, because the estimated coefficients are not invariant but will change along with agents' decision rules in response to a new policy. A classic example of this fallacy was the erroneous inference that a regression of inflation on unemployment (the Phillips curve) represented a structural trade-off for policy to exploit.

Thus, the argument of Lucas (1976) is that *future* regression coefficients and decision rules will be different from those estimated *presently* if the government policy rule changes in the future. Our argument does not concern changes in future regression coefficients. Rather, we argue there will be *a change in what is measured presently* (e.g. the measured control-treatment difference in an RCT) in light of expectations regarding whether and how econometric evidence, causal parameter estimates, will be used in subsequent policy decisions. The second key difference from Lucas is that he considers an (utterly unexpected) *exogenous* change in policy. In contrast, central to our arguments is that evidence-based policy changes are (by definition) endogenous. It is the *endogeneity* of the policy change that gives rise to the five novel econometric challenges we demonstrate, especially the role of the government objective function, the role of prior beliefs regarding parameters to be estimated, and the confounding effects of endogenously heterogeneous beliefs. Third, and finally, in the argument of Lucas (1976), econometricians sit outside the model in that their estimates are not part of the information set of agents inside the model. In contrast, econometricians sit inside our model, with our focus being on the feedback between econometricians' estimates and government policy. This feedback is an underlying cause of the novel biases and paradox we illustrate, upon which Lucas (1976) is silent.

The macro-econometric literature has focused on the implications of rational expectations for the interpretation of vector autoregressions. Sargent (1971, 1973, 1977) and Taylor (1979) showed that rational expectations implies restrictions on distributed lags of vector autoregressions. Sims (1982) and Sargent (1984) pointed to an asymmetry in rational expectations econometrics practice in

postulating optimizing behavior on the part of households and firms while assuming non-optimizing behavior by the government. In our paper, we discuss the correct interpretation of experimental data assuming all agents, including the government, behave optimally and make optimal use of the information available to them.

This paper is also related to that of Hennessy and Strebulaev (2015) who analyze the meaning of econometric evidence derived from an economy hard-wired with an infinite sequence of exogenous natural policy experiments, with zero endogeneity bias at any stage. In contrast, we here consider an economy with only two policy changes. The first policy change arises from an exogenous natural experiment. The second policy change is an optimal response to evidence derived from the first. It is this second-stage governmental policy optimization that is the source of the biases and paradox we discuss.

Our critique of the random assignment literature is related to, but distinct from the critique made by Heckman (1997) who argues that agents can be expected to endogenously undermine random assignment when beneficial to do so. In our laboratory economies, firms are incapable of avoiding the experimental policy treatment. Heckman (1997) and Deaton (2010) emphasize that with heterogeneity the probability limit of instrumental variables estimators can depend on the choice of instrument. In our model, there is no instrumentation. Deaton also emphasizes practical real-world problems associated with small samples and bias in the selection of sample panels. We consider infinite sample sizes, in that there is a continuum of treated and control firms, with ideal randomization.

A Hawthorne Effect is said to arise if treated agents change their behavior under observation (see Levitt and List (2011) and Zwane et al. (2011)). A John Henry Effect is said to arise if control group agents change their behavior under observation. These observer effects are commonly discussed in the behavioral and organization literatures. These literatures have postulated a range of behavioral rationales for such effects such as self-consciousness, approval-seeking, spite, or a desire to influence study outcomes. Importantly, our model abstracts from such effects since firms are rational, measure zero, and anonymous.

Acemoglu (2010) argues that general equilibrium effects can undermine the external validity of (small-scale) experimental evidence, with large-scale policy changes potentially leading to factor substitution and/or endogenous changes in prices and technology. All such effects are shut off in our model. Acemoglu also points out that inherent differences in technology or institutions can limit external validity. These effects are also shut off in our model as we consider inference within a single parable economy.

The rest of the paper is as follows. Section 2 provides some recent examples of evidence-based policymaking. Section 3 presents a model of the interaction between firms, governments, and econo-

metricians. Section 4 uses the model as the basis for discussing econometric inference in settings where firms face a common exogenous policy shock (NPEs). Section 5 discusses observer effects in NPEs. Section 6 discusses RCTs.

2 Evidence-Based Policy-Making: Examples

To illustrate the relevance of our analysis, consider first the area of environmental economics. In conjunction with President Obama's landmark Clean Policy Plan, the U.S. Environmental Protection Agency was required to submit a detailed regulatory impact analysis. Reflecting the gold standard status accorded to evidence from quasi-randomizations, the EPA's analysis of employment impacts relies heavily on evidence from quasi-natural experiments. For example, the EPA Report (2014) cites Greenstone (2002) who assessed the economic impact of the Clean Air Amendments based on variation in regulation resulting from quasi-random changes in pollution attainment status across U.S. counties over time. Other papers relying on this same source of quasi-random variation, and singled out in the EPA's report, include Berman and Bui (2001), Walker (2011), and Kahn and Mansur (2013). Walker (2011) highlights the evidence-policy loop (2011) in writing, "My estimates from the most recent revisions are arguably more applicable to current policy debates, and are particularly important in light of the EPA's recent proposal to further strengthen emissions standards."

In the area of development economics, the type of randomized evidence-based policymaking we consider is common. For example, as described by Duflo, Glennerster and Kremer (2006), a prototypical J-PAL study consists in working in close partnership with a sponsoring institution to conduct a pilot program RCT, followed by a broader application of the policy to the treatment, control and other groups during a "scale-up" stage if the pilot was deemed successful. Writing for J-PAL, Dhaliwal and Tulloch (2015) describe close partnerships between experimenters and those setting long-term policy as creating a "virtuous feedback loop." As a specific example of feedback, in describing Mexico's PROGRESA program, they write, "The strength of the evidence coming from these evaluations, as well as the immense popular support the program enjoyed, likely contributed to make it politically infeasible to discontinue the program, and it continues under the new name Opportunidades."

The U.S. Securities and Exchange Commission (SEC) has begun to rely heavily on evidence from random assignment in the policy-setting process. For example, in 2005 and 2006 the SEC temporarily removed the restriction of short-selling to up-ticks on a matched sample of one-third of the Russell 3000 stocks. The evidence was analyzed internally by the Office of Economic Analysis staff on the part of the SEC, with the data also being released publicly facilitating academic study,

e.g. Boehmer, Jones and Zhang (2008). Based on the evidence, the SEC repealed the up-tick rule on all stocks with some pull-back in financial stocks during the 2008 financial crisis. Describing the experiment and feedback effects, the SEC's former Chief Economist Chester Spatt (2011) writes, "I view the very limited nature of the eventual pull-back on what had become such a politically sensitive rule as a reflection of the strength of the original evidence that the SEC staff generated and upon which the repeal had been based."

In the area of health economics, evidence from random assignment has also taken pride of place in the policy arena. Here two studies have been most influential, the RAND Health Insurance Experiment and the Oregon Health Study. Running from 1974 to 1981, the RAND experiment randomly assigned households to health insurance plans with varying levels of cost sharing, allowing for estimation of health and expenditure impacts. Aron-Dine, Einav and Finkelstein (2013) describe feedback effects in writing, "More than three decades later, the RAND results are still widely held to be the gold standard of evidence for predicting the likely impact of health insurance reforms... such estimates have enormous influence as federal and state policymakers consider potential policy interventions to reduce spending on health care."

The Oregon Health Study involves expanded Medicaid access determined via lottery conducted by the state's government in 2008. Allen, Baicker, Finkelstein, Taubman and Wright (2014) link the Oregon Study with the intense debate over President Obama's health reforms, writing:

> One of the primary components of the recently enacted health reform law, the Patient Protection and Affordable Care Act of 2010, is a major expansion of Medicaid, particularly to low income adults. The probable impact of such an expansion on the newly covered population is of obvious interest. This article describes an ongoing Medicaid expansion experiment in Oregon that provides a unique opportunity to investigate its impact through randomized evaluation.

As another example of potential feedback from econometric estimation to policy-setting, consider that during her term as Chairman of the Council of Economic Advisers (2009-2010), Christina Romer was involved in intense debates regarding the likely impact of proposed fiscal stimulus. At the same time, Romer and Romer (2010) published their influential paper assessing the impact of tax changes on economic activity, a study relying on exogenous tax changes identified from legislative narratives. It is not implausible to believe that Romer, in her policy-making role, was heavily influenced by her own econometric evidence, and discounted evidence coming from less credible sources. Romer and Romer (2010) write, "There is pervasive omitted variable bias in any regression of output on an aggregate measure of tax changes... This paper suggests one way of dealing with this omitted variable bias."

Finally, ex post endogeneity will tend to cloud the interpretation of experiments in long-term contracting settings. For example, Khan, Khwaja and Olken (2016) analyze the randomized assignment of property tax assessors in Pakistan to alternative compensation schemes. Clearly, the behavior of governmental officials during the study will be influenced by their expectations regarding the outcome of the experiment in terms of determining the choice of contract post-experiment. For example, the severity of the ratchet-effect depends upon expectations regarding the future performance target, itself likely shaped by the experiment outcomes. Further, expectations regarding the relative weight placed on discretionary payments in the future will influence the willingness of tax collectors to accumulate reputational capital through honest dealing during the experiment.

3 The Model

We begin by contrasting inference in two economies endowed with identical natural experiments and technologies but differing in how evidence is used. In the *Endogenous Policy Economy*, empirical evidence is used to select an optimal long-term policy. In the *Exogenous Policy Economy*, evidence is irrelevant because the government is powerless to change the policy variable.

The underlying investment model is deliberately simple, following, say, Dixit and Pindyck (1994). However, understanding the econometric issues requires a careful articulation of the respective information sets of all agents.

3.1 Technology

Time is continuous and the horizon infinite. All agents are risk-neutral and share a common discount rate $r > 0$. There is a measure one continuum of anonymous firms. The fact that firms are atomistic ensures that, by construction, no firm has any incentive to change its behavior with the goal of influencing econometric inference, and with it, government policy.

We describe the decision problem of an arbitrary firm, omitting time and firm identifiers where obvious to conserve notation. The law of motion for a firm's capital stock is:

$$dk_t = (i_t - \delta k_t)dt. \tag{1}$$

The variable i denotes gross investment and $\delta \geq 0$ is the depreciation rate.

The literature on dynamic accumulation problems, e.g. Abel and Eberly (1994, 1997), has considered complexities arising from fixed costs and irreversibilities. Since such complexities are extraneous to our main arguments, we start by assuming firms face the following convex investment cost function with the goal of keeping the algebra simple:

$$\psi(i) \equiv \gamma i^{\nu/(\nu-1)}. \tag{2}$$

We assume the cost function parameters satisfy $\gamma > 0$ and $\nu > 1$. Imposing $\nu > 1$ ensures convexity of the investment cost function and unique interior optimal policies. As shown below, in this setting one obtains simple closed-form expressions for the empirical outcome variable i for arbitrary $\nu > 1$. The firm's value function is of only peripheral interest, so we relegate solutions for the value function to the appendix, confining attention there to integer values of ν, as in Abel and Eberly (1997).

Firm cash-flow at time t is given by:

$$\Omega(k_t, x_t, \pi_t, i_t, b) \equiv (x_t + \pi_t b)k_t - \gamma i_t^{\nu/(\nu-1)}. \tag{3}$$

The profit factor x in the cash-flow equation is a positive geometric Brownian motion with the following law of motion:

$$dx_t = \mu x_t dt + \sigma x_t dw_t. \tag{4}$$

The variable w denotes an independent Wiener process. To ensure bounded firm value, one must assume the discount rate satisfies $r > \mu\nu + \frac{1}{2}\sigma^2\nu(\nu - 1)$.

We shall think of firms as facing the same x process. This assumption is not necessary, but serves the expositional purpose of generating the type of shocks to which a real-world government might be expected to respond. Analogous sources of ex ante endogeneity bias are generally of paramount concern to econometricians, garnering enormous attention. Anticipating, this and other standard forms of ex ante endogeneity bias will be ruled out by construction so that the experiments considered are optically ideal.

The term $\pi_t b$ entering cash-flow allows us to capture government policy in a very simple manner. The variable π represents a government-stipulated pollution cap. The variable $b \geq 0$ measures an unobservable private benefit to polluting, e.g. cost/effort savings. The benefit to polluting is specific to the firm's industry. Each firm operates in one of $M \geq 2$ industries. Econometricians know which industry each firm falls into, but do not know the marginal benefit, call it b_m, firms in industry m derive from a higher pollution cap. More generally, we could equally well assume that each firm falls into one of a large yet finite number of categories, e.g. the category of textile manufacturers in Kentucky with operating assets above the industry median, say. The key assumption is that econometricians can observe the firm's category, with the category determining the pollution benefit.

The industry-specific pollution benefits are determined as i.i.d. draws at date 0. The pollution benefit for each industry is drawn from the interval $[0, \overline{b}]$ with a strictly positive probability density f on this support, having a corresponding cumulative distribution F that is twice continuously differentiable, with $F(0) = 0$. Tildes denote random variables and bold-type denotes vectors. The realization of the random vector $\widetilde{\mathbf{b}}$, is denoted $\mathbf{b}$. Econometricians and the government will want to infer the vector of industry-level pollution benefit parameters $\mathbf{b}$ in order to inform decisions

regarding the pollution cap post-experiment. In this context, we shall speak of F as capturing prior beliefs.

3.2 Timing

Firms invest optimally each instant. It is convenient to split the model into three stages, $S \in \{P, E, I\}$. Maximum allowed pollution in stage S, denoted π_S, is either $\underline{\pi}$ or $\overline{\pi}$. Stage P is the Pre-Experiment Stage. It is followed by Stage E, the Experiment Stage, which is followed by Stage I, the Implementation Stage. As we show, within each stage firms face a simple time-homogeneous investment problem.

During Stage P, all firms face the same economy-wide pollution cap π_P. This can be thought of as the initial endowed technology in the economy. An exogenous natural experiment will arrive at date τ_E. This date is an independent random variable. The transition rate into the Experiment Stage is $\lambda_E > 0$. It follows that at any time prior to the transition, the expected remaining duration of Stage P is λ_E^{-1}.

During the Experiment Stage, Nature will randomly assign a fraction θ of firms to deregulated status, with $\pi_E = \overline{\pi}$, while remaining firms are assigned to regulated status, with $\pi_E = \underline{\pi}$. Two types of experiments are considered. In a *Natural Policy Experiment* (NPE below), all firms face the same exogenous regulatory policy $\pi_E \neq \pi_P$ during the Experiment Stage. Thus, in our NPEs, the model parameter θ will be set to either 0 or 1. In a *Randomized Controlled Trial* (RCT below), treatment and control groups face different regulations during the experiment. Thus, to capture an RCT we will choose the parameter $\theta \in (0, 1)$. For example, in our RCTs featuring an equal measure of firms in treatment and control groups, the parameter θ will be set to $1/2$.

"Unexpected" experiments are accommodated by setting λ_E to an arbitrarily small number. This would have no effect on treatment and control group differences in an RCT. For NPEs, unexpected experiments feature larger responses at the start of the experiment. This has no effect on the substance of our arguments, as shown below.

The Implementation Stage I will arrive at date τ_I. This date is an independent random variable given τ_E. The transition rate into the Implementation Stage is $\lambda_I > 0$. Thus, at any instant prior to the transition, the expected remaining duration of Stage E is λ_I^{-1}.

At the very start of Stage I, econometricians have the opportunity to observe some empirical evidence.[1][2] In the RCT, econometricians can look back and measure the difference between the investment of treatment and control groups, industry by industry, at an arbitrary point in time during the experiment. In the NPE, econometricians can look back and measure the change in

[1] Observation during Stage E adds an uninteresting limbo phase where π_I is inferred, with π_E still in effect.

[2] Letting firms make the same observation as econometricians has no effect.

industry-level investment that took place at the start of the experiment. Since time is continuous and x path-continuous, measuring investment changes over the instants just before and just after experiment initiation obviates the need to control for changes in economic conditions (the state variable x in our model). As shown below, the parameter vector of interest here $\mathbf{b}$, can generally be correctly inferred from these standard empirical statistics, but only if the econometrician understands the interplay between evidence, policy, and firm-level expectations.

The following assumption is satisfied by the stochastic policy process facing each firm j at any point in time at which the experimental measurement may take place.

Assumption 1 (Independence): $\widetilde{\pi}_{jt} \perp \{\widetilde{\mathbf{b}}, \widetilde{x}_t\} \quad \forall \quad j$ and $t \in [0, \tau_I)$.

By construction, independence of the policy process rules out the standard forms of endogeneity bias about which empiricists would be expected to debate almost exclusively in a real-world seminar. First, random assignment rules out selection by firms or the government based on unobservables ($\mathbf{b}$), e.g. heavy polluters choosing jurisdictions less likely to face experimental regulation. Second, econometricians might also be concerned that the government they observe is only willing to impose the novel regulation (π_E) because it knows that the true parameter configuration is benign, or has favorable priors. However, in the economy considered, the government does not know the parameter vector $\mathbf{b}$, and the policy variable is independent of the random vector $\widetilde{\mathbf{b}}$. Finally, we recall that in equation (4) the Wiener process w_t was assumed to be independent, eliminating concern that experimental regulation is correlated with the underlying profitability state (x_t).

Since we will be evaluating the difference-in-difference estimator, it is worth noting that our experiments will have very good optics. For example, it follows from Assumption 1 (Independence) and Ito's lemma that, within each industry, the investment of treatment and control groups will exhibit identical stochastic trends (and levels, actually) prior to the experiment.[3] In particular, for all firms j within a given industry m we have:

$$\textbf{Common Trends: } di_{jt} = \left[\mu x_t i_x^P(x_t, b_m) + \frac{1}{2} i_{xx}^P(x_t, b_m)\sigma^2 x_t^2\right] dt + \frac{1}{2}\sigma^2 x_t^2 i_x^P(x_t, b_m) dw_t \quad \forall \quad j \text{ and } t < \tau_E. \tag{5}$$

The Pre-Experiment Stage investment policy function i^P has an analytical solution derived below.

During the Implementation Stage, a long-term regulatory policy will be implemented permanently. In the Exogenous Policy Economy, π_I will be set to some known technologically predetermined value π_I^{EX} and the government is powerless to alter this fact. In the Endogenous Policy Economy, the government will, with the help of its sophisticated econometrician, infer the parameter vector $\mathbf{b}$ and then implement an optimal regulatory policy in light of this information.

[3] In fact, the trend is common across industries given the investment function.

3.3 Firm-Level Policy Beliefs

In the Endogenous Policy Economy, once the Implementation Stage is reached, the government will put into place an optimal policy according to its objective function Θ. The objective function is common knowledge to all agents at date 0. To keep the math simple, we'll assume the government will set policy based on the average industry-level pollution benefit. We'll consider one of two possibilities. The government can be Pro-Environment or it can be Pro-Investment. Respectively, such governments have the following objective functions.

$$\textbf{Pro-Environment} \quad : \quad \Theta(\pi, \mathbf{b}) \equiv \pi \times \left[b^* - \frac{1}{M} \sum_{m=1}^{M} b_m \right]. \tag{6}$$

$$\textbf{Pro-Investment} \quad : \quad \Theta(\pi, \mathbf{b}) \equiv \pi \times \left[\frac{1}{M} \sum_{m=1}^{M} b_m - b^* \right].$$

In the preceding equations $b^* > 0$ is a cutoff for the average pollution benefit. A Pro-Investment government will deregulate if the average pollution benefit is sufficiently high, since the post-experiment investment of each firm j is increasing in its benefit parameter b_j. A Pro-Environment government will instead regulate if the ability to pollute is sufficiently valuable to the average firm, since deregulation would then lead to a large increase in firm investment in polluting capital stocks, investment this government type wants to deter.

We consider Rational Expectations Equilibria in which the government is able to infer $\mathbf{b}$ based on the econometric evidence. Under endogenous policies, the government will implement:

$$\pi_I^*(\mathbf{b}) \in \arg \max_{\pi \in \{\underline{\pi}, \overline{\pi}\}} \quad \Theta(\pi, \mathbf{b}). \tag{7}$$

We will impose the following assumption which ensures that econometric inference actually serves some real purpose in the Endogenous Policy Economy.

$$\textbf{Evidence Policy-Relevant}: \; b^* < \overline{b}. \tag{8}$$

Notice, if this condition were not satisfied then the government would never change its policy regardless of the econometric evidence. The Pro-Investment government would never deregulate and the Pro-Environment government would always deregulate.

We can then define an indicator for the optimality of deregulation given the (correctly) inferred parameter vector:

$$\overline{\pi} \in \arg \max_{\pi \in \{\underline{\pi}, \overline{\pi}\}} \quad \Theta(\pi, \mathbf{b}) \Leftrightarrow \chi(\mathbf{b}) = 1. \tag{9}$$

The policy beliefs for a firm in industry m are summarized by the following belief-function which measures their assessment of the probability of deregulation being implemented post-experiment:

$$\beta(b) \equiv \int_0^{\bar{b}} ... \int_0^{\bar{b}} \int_0^{\bar{b}} ... \int_0^{\bar{b}} \left[\begin{array}{c} [\chi(b_1, ..., b_{m-1}, b, b_{m+1}, ..., b_M)] \\ [F(db_1)...F(db_{m-1})F(db_{m+1})...F(db_M) \end{array} \right]. \tag{10}$$

In the case of a Pro-Investment (Pro-Environment) government, it is apparent that the indicator function in the preceding equation is non-decreasing (non-increasing) in the own-industry benefit parameter b. Further, from equation (8) it follows that the final government policy rule is always uncertain and that the inferred parameter for each industry influences marginal decisions. We thus have the following lemma.

Lemma 1 *If the government is Pro-Investment (Pro-Environment) and free to implement its optimal policy, each firm j's assessed probability of deregulation post-experiment, or belief function, is weakly increasing (decreasing) in its own causal parameter b_j, and strictly so on a set of positive measure. Firms have heterogeneous post-experiment policy beliefs on a set of positive measure.*

To close this subsection, it will be useful to consider the simplest possible case of two industries. Here the belief function takes the form:

$$\begin{aligned} \textbf{Pro-Environment} \quad &: \quad \beta(b) = F(2b^* - b) \Rightarrow \beta'(b) = -f(2b^* - b) \leq 0 \\ \textbf{Pro-Investment} \quad &: \quad \beta(b) = 1 - F(2b^* - b) \Rightarrow \beta'(b) = f(2b^* - b) \geq 0. \end{aligned} \tag{11}$$

It is apparent from the preceding equations that, consistent with the preceding lemma, endowed heterogeneity in the realized treatment response parameters (b_1, b_2) results in endogenous heterogeneity in firm-level expectations regarding government policy decisions post-experiment. Further, it is apparent that the shape of the belief function is dictated by prior beliefs regarding the parameters to be estimated (F), as well as the government objective function parameter b^*.

These effects are illustrated in Figure 1 which plots realized beliefs as a function of the realized own-industry pollution benefit parameter. The figure assumes F is the uniform distribution on $[0, 1]$ and plots beliefs under both Pro-Environment and Pro-Investment Governments. Consistent with the preceding lemma, beliefs are decreasing under the former and increasing under the latter. The figure considers governments applying low ($b^* = .75$) and high ($b^* = .85$) values for the cutoff parameter. Apparently, changes in this parameter of the government objective lead to endogenous shifts in belief functions, consistent with equation (11).

It is worth closing this subsection by noting that it would be incorrect to assert that "we simply have to make the right assumption about the probability of deregulation." To see this, note first

that there is no common policy expectation. Second, note that making correct assumptions about each agents' policy expectation, a priori, is tantamount to assuming the econometrician knows the very information (**b**) that the econometric estimation is intended to recover.

Finally, while the modeling details might differ depending on context, we would argue that the heterogeneous expectations effect derived here might well be expected to be ubiquitous in settings with evidence-based policymaking: A measure zero firm benefits (suffers) a great deal from some experimental policy, but knows its technology is correlated with those of other firms, so expects that the government, capable of correct statistical inference, is more likely to continue or scale-up (discontinue or scale-down) the policy. As we show, the failure to account for this endogenous heterogeneous beliefs channel leads to incorrect inference regarding causal parameters.

3.4 Investment Decisions

A formal solution using optimal control is provided in the appendix, with analytical solutions derived for integer values of ν. This subsection confines attention to characterizing the empirical outcome variable, firm investment. Despite considering an infinite horizon and continuous-time, the model solution boils down to three stages and can be solved as readily as a three period model.

We begin by focusing on the Endogenous Policy Economy. At each instant it is optimal to invest up to the point that the shadow value of a unit of installed capital, call it q, is just equal to the marginal investment cost. We can therefore define the following function i^* mapping the shadow value to optimal investment:

$$q = \psi'(i^*) \Rightarrow i^*(q) \equiv \left(\frac{\nu - 1}{\nu\gamma}\right)^{\nu-1} q^{\nu-1}. \tag{12}$$

Optimal investment is increasing in q, with the q-sensitivity of investment varying with the parameters of the investment cost function, γ and ν.

Accounting for the fact that capital depreciates at rate δ, the discounted value derived from one unit of capital at its installation date, call this date τ, is given by:[4]

$$\begin{aligned} q(x_\tau, b) &= \int_0^\infty e^{-rt} E_\tau \left\{ e^{-\delta t} \Omega_k(k_{\tau+t}, x_{\tau+t}, \pi_{\tau+t}, i_{\tau+t}, b) \right\} dt \\ &= \int_0^\infty e^{-(r+\delta)t} E_\tau \left\{ x_{\tau+t} + b\pi_{\tau+t} \right\} dt. \end{aligned} \tag{13}$$

Installed capital is valued by applying an effective discount rate of $r + \delta$ to the marginal product from a unit of capital which, in turn, is equal to the profit factor x plus the pollution benefit $b\pi$.

[4] For a derivation, see the appendix or Abel and Eberly (1994, 1997).

Using the Gordon Growth Formula to value the claim to x, which grows at rate μ, it follows that the shadow value of capital at each stage $S \in \{P, E, I\}$ takes the form:

$$q^S(x,b) = \frac{x}{r+\delta-\mu} + \eta^S(b). \tag{14}$$

Above, $\eta^S(b)$ represents the present value future pollution benefits, which can be derived using backward-induction.

During the Implementation Stage, the pollution benefit represents a constant perpetuity. Applying the effective discount rate $r+\delta$, we have:

$$\eta^I(b) = \frac{\pi_I b}{r+\delta}. \tag{15}$$

Throughout, we let upper (lower) bars denote values and policies if during the current stage $\pi = \overline{\pi}$ ($\pi = \underline{\pi}$). We have:

$$\begin{aligned} \overline{q}^I(x,b) &= \frac{x}{r+\delta-\mu} + \frac{\overline{\pi} b}{r+\delta} \\ \underline{q}^I(x,b) &= \frac{x}{r+\delta-\mu} + \frac{\underline{\pi} b}{r+\delta} \end{aligned} \tag{16}$$

Consider next the value of future pollution benefits evaluated during the Experiment Stage. The expected rate of return must equal the effective discount rate $r+\delta$. The return here consists of the flow of pollution benefits plus the capital gain accruing if there is a transition to the Implementation Stage. We have the following equilibrium condition:

$$\begin{aligned} (r+\delta)\eta^E(b) &= \pi_E b + \lambda_I \left[\beta(b)\overline{\eta}^I(b) + (1-\beta(b))\underline{\eta}^I(b) - \eta^E(b)\right] \\ &\Rightarrow \quad \eta^E(b) = \left[\frac{(r+\delta)\pi_E + \lambda_I[\beta(b)\overline{\pi} + (1-\beta(b))\underline{\pi}]}{(r+\delta)(r+\delta+\lambda_I)}\right] b. \end{aligned} \tag{17}$$

Therefore, the shadow value of capital during the Experiment Stage is:

$$\begin{aligned} \overline{q}^E(x,b) &= \frac{x}{r+\delta-\mu} + \left[\overline{\pi} + \lambda_I\left(\frac{\beta(b)\overline{\pi} + (1-\beta(b))\underline{\pi}}{r+\delta}\right)\right]\left[\frac{b}{r+\delta+\lambda_I}\right] \\ \underline{q}^E(x,b) &= \frac{x}{r+\delta-\mu} + \left[\underline{\pi} + \lambda_I\left(\frac{\beta(b)\overline{\pi} + (1-\beta(b))\underline{\pi}}{r+\delta}\right)\right]\left[\frac{b}{r+\delta+\lambda_I}\right]. \end{aligned} \tag{18}$$

Consider finally the value of future pollution benefits evaluated during the Pre-Experiment Stage. The return here consists of the flow of pollution benefits plus the capital gain accruing if there is a transition to the Experiment Stage. Accounting for the fact that the firm faces $\pi_E = \overline{\pi}$ with probability θ, we have the following equilibrium condition:

$$\begin{aligned} (r+\delta)\eta^P(b) &= \pi_P b + \lambda_E \left[\theta\overline{\eta}^E(b) + (1-\theta)\underline{\eta}^E(b) - \eta^P(b)\right] \\ &\Rightarrow \quad \eta^P(b) = \left[\frac{\pi_P}{r+\delta+\lambda_E} + \lambda_E\left(\frac{(r+\delta)(\theta\overline{\pi} + (1-\theta)\underline{\pi}) + \lambda_I\left(\beta(b)\overline{\pi} + (1-\beta(b))\underline{\pi}\right)}{(r+\delta)(r+\delta+\lambda_I)(r+\delta+\lambda_E)}\right)\right] b. \end{aligned} \tag{19}$$

Therefore, the shadow value of capital during the Pre-Experiment Stage is:

$$\overline{q}^P(x,b) = \frac{x}{r+\delta-\mu} + \left[\frac{\overline{\pi}}{r+\delta+\lambda_E} + \lambda_E\left(\frac{(r+\delta)(\theta\overline{\pi}+(1-\theta)\underline{\pi}) + \lambda_I\left(\beta(b)\overline{\pi}+(1-\beta(b))\underline{\pi}\right)}{(r+\delta)(r+\delta+\lambda_I)(r+\delta+\lambda_E)}\right)\right] b \quad (20)$$

$$\underline{q}^P(x,b) = \frac{x}{r+\delta-\mu} + \left[\frac{\underline{\pi}}{r+\delta+\lambda_E} + \lambda_E\left(\frac{(r+\delta)(\theta\overline{\pi}+(1-\theta)\underline{\pi}) + \lambda_I\left(\beta(b)\overline{\pi}+(1-\beta(b))\underline{\pi}\right)}{(r+\delta)(r+\delta+\lambda_I)(r+\delta+\lambda_E)}\right)\right] b.$$

Consider next the Exogenous Policy Economy. The only necessary modification to the preceding analysis is that here beliefs regarding π_I must coincide with the known pre-determined value π_I^{EX}. Thus, one must simply introduce the following substitution into the Endogenous Policy Economy shadow value solutions:

$$\beta(b)\overline{\pi} + (1-\beta(b))\underline{\pi} \to \pi_I^{EX}. \quad (21)$$

In other words, in the Exogenous Policy Economy beliefs are homogeneous, invariant to the own-industry pollution benefit, and extreme in that the belief β is either 0 or 1.

Finally, before closing this subsection it is worth pinning down the causal effect (CE) parameters in this economy. As argued by Heckman (2000), a causal effect is just a Marshallian comparative static. Recall, in our economy, the government is interested in predicting how each industry's investment will vary with the pollution cap it puts into place at the start of the Implementation Stage. For firms in industry m, the causal effect is:

$$CE_m \equiv \frac{\partial i^I(x,b_m;\pi)}{\partial \pi} = \frac{\partial q^I(x,b_m;\pi)}{\partial \pi}\frac{di^*}{dq} = \left(\frac{b_m}{r+\delta}\right)(\nu-1)\left(\frac{\nu-1}{\nu\gamma}\right)^{\nu-1}[q^I(x,b_m;\pi)]^{\nu-2}. \quad (22)$$

It follows from the preceding equation that in order for the government to correctly predict the response to its policy decision, it must correctly infer the vector $\mathbf{b}$ of causal parameters.

4 Inference in Natural Policy Experiments

This section considers the following NPE. During the Pre-Experiment Stage, there is regulation, with $\pi_P = \underline{\pi}$. During the Experiment Stage all firms face deregulation, with $\pi_E = \overline{\pi}$. In the Exogenous Policy Economy, the policy variable remains fixed at $\pi_I^{EX} = \overline{\pi}$ after the experiment and the government is powerless to alter this fact. In the Endogenous Policy Economy, the government's econometrician correctly infers the vector $\mathbf{b}$ based on industry-specific responses. The government then implements its optimal policy. In this setting, we evaluate inference by an econometrician sitting outside the government, say an academic economist.

For the purpose of the numerical examples, we assume there are two industries, with their pollution benefits being i.i.d. draws from the uniform distribution on $[0,1]$. We begin by considering the Endogenous Policy Economy under a Pro-Environment government that will deregulate only if

the average pollution benefit falls below the cutoff value b^* =0.60. We'll consider here relatively long expected policy shock durations of 6.7 years, as would tend to be true for economy-wide legislation. The following parameter values are assumed: $r = .05$; $\delta = .10$; $\mu = 0$; $\lambda_E = .15$; $\lambda_I = .15$; $\overline{\pi} = 1$; $\underline{\pi} = 0$; $x = 1$; $\gamma = 1$; and $\nu = 2$. It is worth noting that the numerical illustrations assume adjustment costs are quadratic ($\nu = 2$) which implies optimal investment is linear in q. Although this parametric assumption is restrictive, it allows us to most clearly illustrate the role of belief heterogeneity in NPEs.

Consider now Figure 2. On the horizontal axis is the true pollution benefit for one of the two industries, a parameter our econometrician wants to infer. On the vertical axis is the change in investment at the start of the experiment. Consider the contrast between the response in the Endogenous versus Exogenous policy economies. In both cases, there is zero investment response if $b = 0$. After all, if profits are invariant to the pollution cap, investment will be unresponsive to the cap. Further, we see the investment response is monotonically increasing in b. Since here the measured response is monotonic in the parameter, the econometrician will be able to correctly infer its value provided she accounts for the interplay between estimation, expectations, and policy-setting.

Suppose our academic econometrician works in the Endogenous Policy Economy. If the econometrician is sophisticated, she will anticipate the future utilization of econometric evidence as a policy input and use the solid curve in performing her own inference, resulting in correct estimation of b. If she is naïve, the policy input role is ignored and she instead uses the dashed line to perform inference. It is apparent the naïve econometrician will understate pollution benefits. For example, if the realized $b = 1$, the observed investment response will be 1. However, working along the dashed line the naïve econometrician will infer the pollution benefit is only $b = .60$.

Intuitively, in the Endogenous Policy Economy, the government has the power to reverse Nature's course and re-impose regulation based on the experimental data. Consequently, the positive investment response to experimental deregulation will be less dramatic. This is an intuitive effect. But note, the wedge resulting from this expectations effect is not homogeneous. Rather, it is apparent that the wedge grows ever larger as b grows larger. As shown formally below, this non-linearity is due to the fact that a higher industry-level pollution benefit causes the firm to rationally infer a lower likelihood that the Pro-Environment government will deregulate long-term, i.e. $\beta' \leq 0$. Heterogeneous causal parameters lead to heterogeneous beliefs, affecting what is measured. And here the failure to account for the compression of treatment responses, across b values, resulting from endogenous belief heterogeneity, apparently leads to an understatement of the true value of the causal parameter.

Another issue of concern to econometricians is that policy endogeneity limits external validity.

For example, if a government happened to know up-front ($t = 0$), that $\mathbf{b}$ in its jurisdiction is atypically low, it would be more willing to introduce the deregulation considered. Assumption 1 rules out this form of ex ante endogeneity. However, ex post endogeneity still leads to similar external validity challenges. For example, instead of happening to know $\mathbf{b}$ is low, suppose instead there is an otherwise equivalent economy having more favorable priors (F), with F now being a triangular distribution placing higher weight on lower b values. As shown in the dotted line, for any given value of the parameter b, the positive investment reaction would actually be stronger in this Positive Priors Economy. Notice, prior beliefs influence measured responses here even though, by construction, the government abandons its priors when setting long-term policy, since it correctly infers the true parameter values. Obviously, a correct stipulation of priors is necessary to correctly map the observed investment reaction back to the true value of the causal parameter.

As a final example of issues that arise here, suppose government policy is endogenous, but with the government applying a lower cutoff of b^* =0.50. The positive investment response at the start of the experiment will be weaker than in an otherwise identical economy with a higher cutoff. Clearly, correct inference of the causal parameter is predicated upon a correct stipulation of the parameters of the government objective function into which the parameter estimates will be fed.

We can easily understand the inference problem here analytically. The measured change in investment at the start of the experiment (the vertical axis in Figure 2), call it the deregulation response function (R), is:

$$\begin{aligned} R(b) &\equiv i^*[\overline{q}^E(x,b)] - i^*[\underline{q}^P(x,b)] \\ &= \left(\frac{\nu-1}{\nu\gamma}\right)^{\nu-1}\left[\left(\overline{q}^E(x,b)\right)^{\nu-1} - \left(\underline{q}^P(x,b)\right)^{\nu-1}\right] \end{aligned} \tag{23}$$

where the shadow values (q) are as shown in equations (18) and (20). Since these shadow values depend upon beliefs, as captured by the presence of the $\beta(b)$ terms, it is apparent that any change in economic environment (e.g. F and b^*) leading to changes in belief functions results in shifts of the regulatory response function.

To illustrate, consider then the special case of quadratic investment costs ($\nu = 2$). Here the deregulation response function is:

$$\begin{aligned} R(b) &= \left(\frac{1}{2\gamma}\right)\left[\frac{(r+\delta)(\overline{\pi}-\underline{\pi}) + \lambda_I[\beta(b)\overline{\pi} + (1-\beta(b))\underline{\pi} - \underline{\pi}]}{(r+\delta+\lambda_I)(r+\delta+\lambda_E)}\right] b \\ \Rightarrow\ R'(b) &= \left(\frac{1}{2\gamma}\right)\left[\begin{array}{c} \frac{(r+\delta)(\overline{\pi}-\underline{\pi})+\lambda_I[\beta(b)\overline{\pi}+(1-\beta(b))\underline{\pi}-\underline{\pi}]}{(r+\delta+\lambda_I)(r+\delta+\lambda_E)} \\ +\beta'(b)\frac{\lambda_I(\overline{\pi}-\underline{\pi})b}{(r+\delta+\lambda_I)(r+\delta+\lambda_E)} \end{array}\right]. \end{aligned}$$

The preceding equations make clear that a correct stipulation of the belief function $\beta(\cdot)$, which essentially captures endogenous belief heterogeneity, is required to correctly infer the true parameter

b based on an observed investment response. In turn, a correct stipulation of the belief function requires a correct stipulation of priors (F) and parameters of the government's objective function (here b^*). It is apparent from the second line that the response function is linear only if beliefs are constant, as was the case in the Exogenous Policy Economy. Finally, it is also apparent from examining the derivative of the response function that the heterogeneous beliefs channel here attenuates (amplifies) investment response heterogeneity across the different possible endowed b values under a Pro-Environment (Pro-Investment) government.

Finally, if one were to compare the regulation response across economies with endogenous versus exogenous policies, the implied difference is:

$$\underset{Endogenous}{R(b)} - \underset{Exogenous}{R(b)} = \left(\frac{\lambda_I}{2\gamma}\right)\left[\frac{\beta(b)\overline{\pi} + (1-\beta(b))\underline{\pi} - \pi_I^{EX}}{(r+\delta+\lambda_I)(r+\delta+\lambda_E)}\right] b. \tag{24}$$

It follows from equation (24) that there is necessarily a wedge between the experiment response functions in the Exogenous Policy and Endogenous Policy economies.

The following proposition summarizes the results from this section.

Proposition 1 *The deregulation response function (R) for Natural Policy Experiments differs according to whether the econometric evidence is relevant, due to endogenous long-term policy, or irrelevant, due to exogenous long-term policy. Across economies with endogenous long-term policies, deregulation response functions are equal if and only if they share a common government objective function post-experiment and common prior beliefs (F) regarding the distribution governing the parameters to be estimated. Endogenous belief heterogeneity amplifies (attenuates) cross-sectional deregulation investment response heterogeneity under a Pro-Investment (Pro-Environment) government.*

5 Observer Effects in Natural Policy Experiments

This section utilizes the model to examine the link between the perceived credibility of evidence from an NPE and the nature of the evidence itself.

Consider two economies with the same technologies as described in the preceding section, with the two governments sharing the same objective function. Further, assume the same sequence of events as in the previous section. However, assume now that both governments have the ability to choose long-term policy, $\pi_I \in \{\underline{\pi}, \overline{\pi}\}$, at their discretion. The two economies differ however in respective governmental perceptions of the credibility of experimental evidence. In Economy NC, the government views such evidence as non-credible. In Economy C, the government views experimental

evidence as credible. Finally, assume that observation and measurement of firm behavior may or may not be feasible, and that all agents know whether or not observation is feasible.

In Economy NC, long-term policy will be invariant to whether or not observation occurs. If observation occurs, the government ignores the evidence and implements the policy optimal under prior beliefs, call it π^*_{IP}. And if no observation occurs, the government has no choice but to rely on its priors and so again implements π^*_{IP}. For example, in our previous example, the Pro-Environment government would deregulate ($\pi^*_{IP} = \overline{\pi}$) given its prior beliefs since the expected benefit is one-half and the cutoff value for deregulation was 0.60. Clearly, firms in Economy NC will base their investment on the same conjectured long-term policy regardless of whether or not observation occurs. There is no Hawthorne Effect.

In contrast, in Economy C, long-term government policy is contingent upon whether or not observation occurs. If no observation occurs, the government has no choice but to rely on priors and so implements $\pi_I = \pi^*_{IP}$. If observation occurs, the government views the econometric evidence as credible and uses it to infer the parameter vector $\mathbf{b}$. Firms then rationally anticipate the implementation of $\pi^*_I(\mathbf{b})$ and thus form own-sector beliefs $\beta(b)$. As shown next, a Hawthorne Effect then arises from the change in the probability distribution of the policy variable resulting from the act of observation.

The Hawthorne Effect in Economy C can be expressed in terms of the shadow value of capital. In Economy C, firms behave just as they do in the Exogenous Policy Economy if they are not being observed treating $\pi^{EX}_I = \pi^*_{IP}$. In contrast, under observation they behave just as in the Endogenous Policy Economy, anticipating implementation of $\pi^*_I(\mathbf{b})$. It follows that in Economy C, the jump in the shadow value of capital, and hence investment, at the onset regulation depends upon whether the firms are being observed or not. In particular, under observation we have the following jump in the shadow value of capital at the inception of the experimental deregulation:

$$\underset{\text{Observed}}{\left[\overline{q}^E(x,b) - \underline{q}^P(x,b)\right]} = \underset{\text{Not Observed}}{\left[\overline{q}^E(x,b) - \underline{q}^P(x,b)\right]} + \underset{\text{Hawthorne Effect}}{\frac{\lambda_I[\beta(b)\overline{\pi} + (1-\beta(b))\underline{\pi} - \pi^*_{IP}]b}{(r+\delta+\lambda_I)(r+\delta+\lambda_E)}}. \tag{25}$$

We thus have the following proposition.

Proposition 2 *If the government views the Natural Policy Experiment as credible (non-credible), the outcome variable, the change in investment during the experimental treatment period, is (not) contaminated by a Hawthorne Effect.*

To illustrate, Figure 3 plots the experimental response function in Economy C, where NPEs are viewed as a credible source of information. The solid line depicts the response function without observation and the dashed line depicts the response function with observation. Without observation,

the government sets long-term policy based upon prior beliefs and deregulates with probability one ($\pi^*_{IP} = \overline{\pi}$). With observation, sophisticated analysis allows the government to infer $\mathbf{b}$ and so it implements a contingent optimal policy, with $\pi_I = \pi^*_I(\mathbf{b})$. With observation, the response to the experimental deregulation is attenuated. Further, response heterogeneity is attenuated by endogenous belief heterogeneity given $\beta' \leq 0$ under the Pro-Environment government.

6 Randomized Controlled Trials

This section considers econometric inference in the context of RCTs satisfying Assumption 1 and thus stripped of standard self-selection and endogeneity concerns. Pre-experiment, the treatment and control groups in the RCT will exhibit common investment levels and identical trends (equation (5)).

6.1 Inference in RCTs

To fix ideas, consider then the following RCT setting. Firms are initially deregulated, with $\pi_P = \overline{\pi}$. During the Experiment Stage, an equal measure ($\theta = 1/2$) of firms are assigned to the two possible π values. In the Exogenous Policy Economy, once the Implementation Stage begins, the government is powerless to prevent the policy variable reverting back to $\pi_I = \overline{\pi}$. In the Endogenous Policy Economy, the government sets policy optimally given the evidence supplied by its econometrician.

For the purpose of the numerical example, we assume there are two industries, with their pollution benefits being i.i.d. draws from the uniform distribution on $[0, 1]$. For balance, we now consider the Endogenous Policy Economy under a Pro-Investment government that will deregulate if the average pollution benefit exceeds the cutoff value b^* =0.75. We'll consider here relatively short policy shock durations of 1 year, as would tend to be true for many real-world RCTs. The following parameter values are assumed: $r = .05$; $\delta = .10$; $\mu = 0$; $\lambda_E = 1$; $\lambda_I = 1$; $\overline{\pi} = 1$; $\underline{\pi} = 0$; $x = 1$; $\gamma = 1$; and $\nu = 4$.

On the horizontal axis of Figure 4 is the true value of the unknown parameter b. On the vertical axis is the difference between investment by the control group ($\pi_E = \overline{\pi}$) and treatment group ($\pi_E = \underline{\pi}$). Since treatment and control group firms, being considered industry by industry, have identical investment prior to the experiment, the vertical axis also measures the difference in differences. The solid line measures the control-treatment investment difference in the Endogenous Policy Economy and the dotted-dashed line measures the difference in the Exogenous Policy Economy. The dashed line considers the effect of more favorable priors (triangular distribution with greater weight on high b values) on an otherwise identical economy with endogenous long-term regulation. The dotted

line considers the effect of a higher cutoff value (b^* =0.85) on an otherwise identical economy with endogenous long-term regulation.

Figure 4 allows us to contrast the inference that will be made by a sophisticated econometrician, who accounts for the role of estimation in policy-setting, versus a naïve econometrician who ignores it. Consider, say, an academic econometrician working in the Endogenous Policy Economy. The solid line in Figure 4 reflects this reality. Therefore, if the econometrician is sophisticated, she will account for the link between policymaking and empirical evidence and use the solid line in performing inference, resulting in correct estimation of the parameter b. If the econometrician is naïve, she ignores the link and instead uses the dotted-dashed line to perform inference. The naïve econometrician will understate the value of the causal parameter b. For example, suppose the true value is $b = 1$, resulting in an observed control-treatment investment difference of 100. Working along the dotted-dashed line the naïve econometrician will infer this difference resulted from $b = .70$.

Another oft-mentioned real-world concern is that inferences regarding regulatory impacts will be biased if predicated upon a discretionary RCT, perhaps due to a government having private knowledge. Assumption 1 rules out this problem. However, ex post endogeneity gives rise to a similar problem. To see this, suppose our academic econometrician examines the control-treatment investment difference in an endogenous policy economy endowed with more favorable priors (F) regarding the parameter b. As shown in the dashed line, for any given realization of b, the investment difference is larger in this economy than under less favorable priors (solid line). If the econometrician failed to account for the effect of more favorable priors, she would overstate b. For example, suppose the true value of b is .90, resulting in an observed control-treatment investment difference of 100 under positive priors. Working along the solid line the econometrician will incorrectly conclude this difference resulted from $b = 1$.

It is also apparent from Figure 4 that correct parameter inference in our RCT is contingent upon a correct stipulation of the deep parameters of the government objective function into which econometric estimates will be fed. To see this, note that the dotted line illustrates that a shift downward in the control-treatment investment difference occurs here if the government were to impose a higher threshold for deregulation.

6.2 Hawthorne and John Henry Effects in RCTs

This subsection considers the potential for control and treatment groups to exhibit observer effects in RCTs. To illustrate, we return to the same RCT and parameter values as in the preceding subsection, focusing on a government that is willing to use evidence from the RCT to set regulatory policy. But again, let us assume that observation may not be feasible, allowing us to assess whether

the act of observation changes behavior.

The results are illustrated in Figure 5. On the horizontal axis is the true value of the causal parameter b, with the figure showing investment by treatment and control groups, as well as the investment difference, for cases when firms are observed and when they are not. As shown, both treatment and control groups change their investment under observation. But note, by construction, the experimental treatment itself is identical under observation versus non-observation: 50% of the firms are regulated and 50% are unregulated.

What does differ between observation and non-observation is the expected path of the policy variable post-experiment. If observed, firms expect the government to utilize the experimental evidence in order to correctly infer $\mathbf{b}$, going on to implement $\pi_I^*(\mathbf{b})$, implying regulation will occur some percentage of the time, with heterogeneous beliefs regarding the probability. Absent observation, firms know the government must rely upon prior beliefs in setting policy long-term, implying regulation with probability one ($\pi_{IP}^* = \underline{\pi}$) given that the assumed value for the deregulation threshold exceeds the unconditional average of b.

Apparently, as shown in Figure 5, changes in the distribution of the policy variable post-experiment, resulting from observation, induce changes in investment by both treatment and control groups during the experimental period, as well as changes in the investment difference. The next subsection sets out to understand why.

6.3 Analytical Treatment of RCTs

This subsection characterizes analytically some underlying challenges to inference in RCTs in dynamic settings. To begin, it will be useful to consider the difference between the shadow value of capital across the control ($\pi_E = \overline{\pi}$) and treatment ($\pi_E = \underline{\pi}$) groups. Using equation (18) we have:

$$\overline{q}^E(x,b) - \underline{q}^E(x,b) = \frac{(\overline{\pi} - \underline{\pi})b}{(r + \delta + \lambda_I)} \tag{26}$$

Notice, the preceding equation shows that the difference between the shadow value of capital between control and treatment groups is actually invariant to the distribution of the policy variable during the Implementation Stage. After all, just as randomized assignment ensures there is no selection based upon unobservable firm characteristics (b_j), random assignment also ensures there is no selection based upon expectations. That is, post-experiment policy expectations are necessarily equalized across treatment and control groups. Since long-term expectations are the same, the difference between the shadow values of capital between treatment and control groups must be attributable to differences in the expected discounted marginal product of capital during the experiment itself. In fact, the difference between $\overline{q}^E$ and $\underline{q}^E$ shown in equation (26) is just the present

value of a claim to the flow of excess benefits $(\overline{\pi}-\underline{\pi})b$ accruing to the deregulated group during the Experiment Stage.

But recall, the preceding subsection found that the control-treatment investment difference did indeed hinge upon expectations regarding the policy variable path post-experiment. It is this investment difference that is the outcome variable observed by econometricians, not the control-treatment shadow value difference (equation (26)). It is the behavior of the former that we must then understand.

Letting Δ denote the difference between control and treatment group investment, from equation (18) we have:

$$\begin{aligned}
\Delta(x,b) &= \underset{Control}{i^*[\overline{q}^E(x,b)]} - \underset{Treatment}{i^*[\underline{q}^E(x,b)]} \qquad (27)\\
&= \left(\frac{\nu-1}{\nu\gamma}\right)^{\nu-1}\left[\left(\underline{q}^E(x,b)+\frac{(\overline{\pi}-\underline{\pi})b}{(r+\delta+\lambda_I)}\right)^{\nu-1}-\left(\underline{q}^E(x,b)\right)^{\nu-1}\right]\\
&= \left(\frac{\nu-1}{\nu\gamma}\right)^{\nu-1}\left[\begin{array}{l}\left[\frac{x}{r+\delta-\mu}+\left(\underline{\pi}+\lambda_I\left(\frac{\beta(b)\overline{\pi}+(1-\beta(b))\underline{\pi}}{r+\delta}\right)\right)\left(\frac{b}{r+\delta+\lambda_I}\right)+\frac{(\overline{\pi}-\underline{\pi})b}{(r+\delta+\lambda_I)}\right]^{\nu-1}\\ \quad -\left[\frac{x}{r+\delta-\mu}+\left(\underline{\pi}+\lambda_I\left(\frac{\beta_H(b)\overline{\pi}+(1-\beta(b))\underline{\pi}}{r+\delta}\right)\right)\left(\frac{b}{r+\delta+\lambda_I}\right)\right]^{\nu-1}\end{array}\right].
\end{aligned}$$

A key point to note in equation (27) is that beliefs ($\beta(b)$) determine the investment of both the treatment and control groups. Since observation influences beliefs, there will be observation effects for both the treatment group (Hawthorne Effect) and the control group (John Henry Effect). Moreover, the size of these effects will vary with prior beliefs and the parameters of the government objective function into which the evidence is fed. After all, as shown in equation (11), the belief function is determined by F and b^*.

Despite the presence of observer effects for both treatment and control groups, it might be hoped that the observer effects will be of equal size across the two groups, so that the control-treatment investment difference will be left uncontaminated. However, as shown in equation (27), in terms of the measured outcome variable i, in contrast to the unmeasured shadow value of capital q, observation effects do not generally cancel.

At this stage it is instructive to consider the exception proving the rule. If and only if one were to assume the investment cost parameter ν is equal to 2, investment is linear in the shadow value of capital and the observation effects hitting both treatment and control groups cancel. In particular, it follows from equation (27) that:

$$\nu=2 \Longrightarrow \Delta(x,b)=\left(\frac{1}{2\gamma}\right)\frac{(\overline{\pi}-\underline{\pi})}{(r+\delta+\lambda_I)}\times b. \qquad (28)$$

Notice, if $\nu=2$, the control-treatment investment difference is linear in the causal parameter b,

and is invariant with respect to expectations regarding the distribution of the policy variable post-experiment.

Up to this point we have considered a restrictive investment cost function to keep the algebra simple. However, in order to provide a complete characterization of the circumstances under which differences and (difference in differences) derived from RCTs are immune from post-experiment expectations contamination, it will be useful to consider a broader class of cost functions, reflective of those considered in the literature with the goal of creating a more realistic depiction of dynamic behavior. A number of realistic frictions create regions of optimal inaction, as well as lumpy policies. For example, Abel and Eberly (1994) consider that there can be fixed costs, and that the agent may not be able to sell capital for the same price at which it is purchased. Chetty (2012) has argued that such frictions and associated inaction regions can cloud the interpretation of empirical evidence. Indeed, as we show next, such frictions contaminate RCTs.

To illustrate, the remaining analysis considers the following General Investment Cost Function.

Definition 1 *General Investment Cost Function: The fixed cost to positive investment is* $\varphi^+ \geq 0$. *The fixed cost to negative investment is* $\varphi^- \geq 0$. *Capital can be purchased at price* P^+ *and sold at price* $P^- \leq P^+$. *Adjustment costs are* ψ, *where* ψ *is a strictly convex twice differentiable function of investment attaining a minimum value of zero at* $i = 0$.

Two points are worth noting at this stage. First, since the General Investment Cost Function shares with the initially-posited cost function ($\gamma i^{\frac{\nu}{\nu-1}}$) the property of being invariant to k, it follows that the shadow value formulae derived in Subsection 3.4 remain valid. Second, Abel and Eberly (1994) show that under such a cost function, investment is weakly monotone increasing in q. They also show that if there are no fixed costs, optimal investment is continuous in q, with $i^* = 0$ optimal for all $q \in [P^-, P^+]$, turning negative at points to the left of this interval and positive at points to the right. If there are fixed costs, optimal accumulation is zero over a wider interval of q values, and exhibits discontinuities at the optimal thresholds for switching from inaction to action.[5]

Recall, under the initially-posited investment cost function, the control-treatment investment difference (as well as the difference in differences) was shown to be invariant to the post-experiment policy variable expectation if and only if investment is linear in q, which held under the parametric assumption $\nu = 2$. To ensure that investment is linear in q under the General Investment Cost Function, one must rule out fixed costs, wedges between the buy and sell price of capital, and assume quadratic adjustment costs. We thus have the following proposition.

[5] See the discussion of Figure 1 in Abel and Eberly (1994).

Proposition 3 *If and only if the Randomized Controlled Trial is relevant, with the empirical evidence affecting the post-experiment policy variable outcome with positive probability, the treatment group will exhibit a Hawthorne Effect and the control group will exhibit a John Henry Effect. The difference between control and treatment group investment (and the difference in their differences) is invariant to factors affecting post-experiment policy variable expectations if and only if the General Investment Cost Function features: a quadratic adjustment cost function (ψ); zero fixed costs ($\varphi^- = \varphi^+ = 0$); and zero wedge between the buy and sell price of capital ($P^- = P^+$).*

The importance of the preceding discussion is illustrated in Figure 6 which plots the control-treatment investment difference as determined by the causal parameter b, while considering alternative configurations of the General Investment Cost Function. Aside from investment costs, the figure retains all the parametric assumptions of Figure 5. Figure 6 now assumes the investment cost parameter ν is equal to 2. The solid line considers the case of non-observation, as well as the case of a cost function meeting the criteria stipulated in the proposition, with zero fixed costs and equality of the buy and sell price of capital, here set at 7.7. Here there is a simple linear relationship between the measured difference and the unknown parameter. Further, here there is no need to account for observation effects when making inferences.

The dashed line considers that there is a wedge between the buy and sell price of capital. In particular, the assumed sell price of capital is only 7, falling below the buy price of capital of 7.7. Clearly, this friction can lead to faulty inference, causing deregulated firms to become inactive for $b \in [.39, .63]$, with regulated firms becoming inactive for $b \in [.60, .74]$. We see that inactivity induces a non-monotonicity in the measured difference. This will complicate inference. For example, if $b \in [.60, .63]$, both groups of firms are inactive and the measured investment difference is zero. If one were to mistakenly rely upon the solid line for inference, ignoring the effect of partial irreversibility, one would conclude that $b = 0$. After all, the investment difference is always equal to zero at $b = 0$.

The dotted line considers that in addition to there being a wedge between the buy and sell price of capital, there is also a small fixed cost (0.20), say a search cost, associated with buying capital. Here one sees that real frictions can create an even more substantial challenge to correct inference. Here the conjunction of real frictions causes deregulated firms to become inactive for $b \in [.39, .77]$, with regulated firms becoming inactive for $b \in [.60, .87]$. We see that inactivity induces two regions of non-monotonicity in the measured difference. This will complicate inference. For example, for $b \in [.60, .77]$, both groups of firms are inactive and the measured investment difference is zero. If one were to mistakenly rely upon the solid line for inference, ignoring the effect of real frictions, one would mistakenly conclude that $b = 0$.

The preceding discussion brings up a more general point. We have been analyzing Rational

Expectations Equilibria in which the government is able to correctly infer causal parameters based on econometric evidence. However, it is apparent that NPEs and RCTs, by themselves, may not be sufficient for this purpose. To see this most clearly, consider first Figure 7A. Figure 7A returns to our initial NPE (Figure 2) featuring deregulation and considers our Pro-Environment government applying an even lower cutoff b^*=.45. We conjecture an equilibrium in which the government can determine the true parameter value. However, suppose the government has only the experimental outcome variable, in our NPE the investment reaction at the start of the experiment, to serve as the basis for its inference. It is apparent that the behavior of this outcome variable is inconsistent with the conjecture of fully correct inference since it is non-monotone in b. The outcome variable simply cannot be inverted to solve for b. Under such a government, we know $\beta' \leq 0$ and so the expectations channel causes industries with higher b values to assess a lower probability of deregulation in the long-term. Apparently, this expectations effect is here strong enough so that there is a region of b values, 0.80 to 0.90, where the investment response is actually decreasing in b. Thus, if the observed investment reaction falls into the region 0.80 to 1, the econometrician would not be able to infer the true b value relying only on the observed investment reaction.

Consider next the feasibility of identification of causal parameters based upon RCTs. In particular, Figure 7B returns to the parameters in Figure 4 but now assumes x is close to zero. As shown, the figure illustrates a case in which a non-monotonicity emerges in the control-treatment investment difference. This non-monotonicity can make it impossible for the government to infer the true value of the causal effect parameter over some ranges of the outcome variable.

7 Conclusion

This paper illustrates a potential tension between the credibility of empirical estimates derived from randomization, and their practical utility. In particular, once this econometric methodology becomes sufficiently credible, and perhaps it already has passed this threshold, estimates derived from it will be used in setting policy. But this can contaminate the original econometric estimates by exposing them to what we have termed ex post endogeneity, with treatment responses dependent upon: the parameters of policymaker objective functions into which estimates are fed; prior beliefs regarding the causal parameters to be estimated; and endogenously heterogeneous policy expectations. As shown, the failure to account for ex post endogeneity leads to faulty inference regarding causal parameters. Further, even with a subtle analysis accounting for the evidence-policy feedback loop, it may not be feasible to infer causal parameters from standard experimental test statistics, with heterogeneous beliefs the key confounding factor. More generally, it is apparent that it is not obvious, a priori, how one should interpret the evidence coming out of RCTs and NPEs in dynamic

settings. Far from being stand-alone objects, correct interpretation may require an extremely subtle analysis or may require the imposition of strong functional form assumptions.

The concerns we raise are most relevant to settings in which agents make forward-looking decisions having expected payoffs that are dependent upon the evidence-based policy decision. Conversely, one can look to the model's key building blocks to identify settings that are less vulnerable to the concerns we raise. First, it is apparent that the importance of our concerns depends upon the proportion of the payoffs to the experiment-phase decision that will accrue post-experiment. For example, if the experiment is very long-lived or the payoffs short-lived, the concerns are less important. Second, the critique is stronger the tighter the nexus between the experiment and the policy decision. It follows then that one may prefer to rely upon other-country evidence or perhaps evidence that is a bit dated, with there then being a relevance versus contamination tradeoff. Finally, it would seem to be optimal that experimental subjects do not understand the link between the experiment and the policy decision. If commitment and policy-discrimination were possible, it might be optimal to commit to not exposing the experimental panel to the long-term optimal policy decision that the experiment informs. Barring such commitment, the experimenter may prefer to keep hidden the policy linkage, but then this would raise ethical concerns.

References

[1] Abel, Andrew, and Janice Eberly, 1994, A unified model of investment under uncertainty, *American Economic Review*.

[2] Abel, Andrew, and Janice Eberly, 1997, An exact solution for the investment and market value of a firm facing uncertainty, adjustment costs, and irreversibility. *Journal of Economic Dynamics and Control* 21, 831-852.

[3] Acemoglu, Daron, 2010, Theory, general equilibrium, and political economy in development economics, *Journal of Economic Perspectives*, 24 (3), 17-32.

[4] Allen, Heidi, Katherine Baicker, Amy Finkelstein, Sarah Taubman, Bill J. Wright, 2010, What the Oregon health study can tell us about expanding medicaid, *Health Affairs* (29), 1498-1506.

[5] Angrist, Joshua D. and Jorn-Steffen Pischke, 2009, *Mostly Harmless Econometrics: An Empiricist's Companion*, Princeton University Press.

[6] Angrist, Joshua D. and Jorn-Steffen Pischke, 2010, The credibility revolution in economics: How better research design is taking the con out of econometrics, *Journal of Economic Perspectives* 24, 3-30.

[7] Aron-Dine, Aviva, Liran Einav, and Amy Finkelstein, 2013, The RAND health insurance experiment, three decades later, *Journal of Economic Perspectives* (27), 197-222.

[8] Berman, Eli and Linda T.M. Bui, 2001, Environmental regulation and labor demand: Evidence from the south coast air basin, *Journal of Public Economics* (79), 265-295.

[9] Boehmer, Ekkehart, Charles M. Jones, and Xiaoyan Zhang, 2008, Unshackling short sellers: The repeal of the uptick rule, Working Paper, Columbia Business School.

[10] Dhaliwal, Iqbal and Caitlin Tulloch, 2015, From research to policy: Evidence from impact evaluations to inform development policy, Working Paper, Abdul Latif Jameel Poverty Action Lab.

[11] Dixit, Avinash and Robert Pindyck, 1994, *Investment Under Uncertainty*, Princeton University Press.

[12] Duflo, Esther, 2004, Scaling up and evaluation, in *Annual World Bank Conference on Development Economics: Accelerating Development*, ed. François Bourguignon and Boris Pleskovic, 341–69. Washington, D.C.: World Bank; Oxford and New York: Oxford University Press.

[13] Greenstone, Michael, 2009, Toward a culture of persistent regulatory experimentation and evaluation, in D. Moss and J. Cisternino, eds. *New Perspectives on Regulation*, Cambridge, MA., The Tobin Project.

[14] Hayashi, Fumio, 1982, Tobin's average and marginal q: A neoclassical interpretation. *Econometrica*.

[15] Heckman, James J., 2000, Causal parameters and policy analysis in economics: A twentieth century retrospective, *Quarterly Journal of Economics* 115 (1), 45-97.

[16] Hennessy, Christopher A. and Ilya A. Strebulaev, 2015, Beyond random assignment: Credible inference of causal effects in dynamic economies, Working paper, London Business School.

[17] Kahn, Matthew E. and Erin T. Mansur, 2013, Do local energy prices and regulation affect the geographic concentration of employment?, *Journal of Public Economics* (101), 105-114.

[18] Khan, Adnan, Asim Khwaja, and Benjamin Olken, 2016, Tax farming redux: Experimenal evidence on performance pay for tax collectors, *Quarterly Journal of Economics* (131), 219-271.

[19] Levitt, Steven D., and John A. List, 2011, Was there really a Hawthorne effect at the Hawthorne plant? An analysis of the original illumination experiments, *American Economic Journal: Applied Economics* 3(1), 224-38.

[20] Ljungqvist, Lars, 2008, Lucas Critique, *The New Palgrave Dictionary of Economics*, Second Edition. Edited by Steven N. Durlauf and Lawrence E. Blume.

[21] Lucas, Robert E., Jr., 1976, Econometric policy evaluation: A critique, in K. Brunner and A. Meltzer Eds., *The Phillips Curve and Labor Markets*, Amsterdam: North Holland.

[22] Romer, Christina D. and David H. Romer, 2010, The macroeconomic effects of tax changes: Estimates based on a new measure of fiscal shocks, *American Economic Review* (100), 763-801.

[23] Sargent, Thomas J., 1971, A note on the accelerationist controversy, *Journal of Money Credit and Banking* 3 (3), 721-725.

[24] Sargent, Thomas J., 1973, Rational expectations and the dynamics of hyperinflation, *International Economic Review* 14 (2), 328-350.

[25] Sargent, Thomas J., 1977, The demand for money during hyperinflations under rational expectations, I, *International Economic Review* 18 (1), 59-82.

[26] Sargent, Thomas J., 1984, Autoregressions, expectations and advice, *American Economic Review* 74, 408-415.

[27] Sims, Christopher A., 1982, Policy analysis with econometric models, *Brookings Papers on Economic Activity* 1, 107-164.

[28] Spatt, Chester, 2011, Measurement and policy formulation, working paper, Carnegie Mellon University, Tepper School of Business.

[29] Taylor, John B., 1979, Estimation and control of a macroeconomic model with rational expectations, *Econometrica* 47, 1267-1286.

[30] Walker, W. Reed, 2011, Environmental regulation and labor reallocation: Evidence from the clean air act, *American Economic Review* (101), 442-447.

[31] Zwane, Alix, Jonathan Zinman, Eric Van Dusen, William Pariente, Clair Null, Edward Miguel, Michael Kremer, Dean S. Karlan, Richard Hornbeck, Xavier Giné, Esther Duflo, Florencia Devoto, Bruno Crepon, and Abhijit Banerjee, 2011, Being surveyed can change later behavior and related parameter estimates, *Proceedings of the National Academy of Science* (February).

Appendix: Model Solution via Optimal Control

For brevity, the argument b is omitted from the derivation, and so the solutions obtained hold for arbitrary b values. In all cases, we pin down analytical solutions for $\nu \in \{2, 3, 4, ...\}$. We solve via backward induction.

Implementation Stage

The Hamilton-Jacobi-Bellman (HJB) equation is:

$$rV^I(x,k) = \max_i (x + \pi_I b)k - \gamma i^{\nu/(\nu-1)} + \mu x V_x^I(x,k) + \frac{1}{2}\sigma^2 x^2 V_{xx}^I(x,k) + (i - \delta k)V_k^I(x,k). \quad (29)$$

We conjecture the following value function that is separable between the value of assets in place and growth options:

$$V^I(x,k) = kq^I(x) + G^I(x). \quad (30)$$

Under the posited functional form, the optimal instantaneous control during the Implementation Stage is $i^*(q^I)$. Substituting the preceding function into the HJB equation, and then isolating the terms scaled by k, we obtain the following ODE for the shadow value of capital:

$$(r + \delta)q^I(x) = (x + \pi_I b) + \mu x q_x^I(x) + \frac{1}{2}\sigma^2 x^2 q_{xx}^I(x). \quad (31)$$

We conjecture the following linear form for the shadow value of capital:

$$q^I(x) = x\zeta^I + \eta^I. \quad (32)$$

Substituting the conjectured solution into the ODE for q we obtain:

$$\begin{aligned} \zeta^I &= \zeta \equiv 1/(r + \delta - \mu). \\ \eta^I &= \frac{\pi_I b}{r + \delta} \end{aligned} \quad (33)$$

This is the shadow value presented in the body of the paper.

We next determine the growth option value function for the Implementation Stage. Substituting the conjectured value function into the HJB equation and dropping now the terms scaled by k that have been eliminated, we obtain the following ODE:

$$rG^I(x) = \mu x G_x^I(x) + \frac{1}{2}\sigma^2 x^2 G_{xx}^I(x) + iq^I(x) - \gamma i^{\nu/(\nu-1)}.$$

We begin by noting that

$$q^I(x) = \zeta x + \eta^I \Rightarrow i^*[q^I(x)]q^I(x) - \gamma[i^*(q^I(x))]^{\nu/(\nu-1)} = (\zeta x + \eta^I)^\nu \Gamma = \sum_{h=0}^{\nu} \phi_h^I x^h \quad (34)$$

where

$$\begin{aligned}\Gamma &\equiv (\nu-1)^{\nu-1}\nu^{-\nu}\gamma^{1-\nu} \\ \phi_h^I &\equiv \binom{\nu}{h}\left(\eta^I\right)^{\nu-h}\zeta^h\Gamma.\end{aligned}$$

The preceding result follows from the binomial expansion formula. Utilizing the binomial expansion result above it follows that the growth option value must satisfy the following ordinary differential equation:

$$rG^I(x) = \mu x G_x^I(x) + \frac{1}{2}\sigma^2 x^2 G_{xx}^I(x) + \sum_{h=0}^{\nu}\phi_h^I x^h.$$

Since the preceding form of growth option value function will recur, it will be convenient to reference the following lemma.

Lemma 2 *The growth option value function satisfying*

$$\widehat{r}G(x) = \mu x G_x(x) + \frac{1}{2}\sigma^2 x^2 G_{xx}(x) + \sum_{h=0}^{\nu}\phi_h x^h$$

has solution

$$\begin{aligned}G(x) &= \sum_{h=0}^{\nu}\phi_h\omega_h x^h \\ \omega_h &\equiv \frac{1}{\widehat{r}-\mu h-\frac{1}{2}\sigma^2 h(h-1)}.\end{aligned}$$

Proof. The function G represents the value of a claim to a sum of geometric Brownian motions to successive powers. The value g_h of a claim to an arbitrary constituent flow payment $\phi_h x^h$ must satisfy the differential equation:

$$\widehat{r}g_h(x) = \mu x g_h'(x) + \frac{1}{2}\sigma^2 x^2 g_h''(x) + \phi_h x^h \tag{35}$$

We conjecture this value function takes the form:

$$g_h(x) = \phi_h\omega_h x^h$$

Substituting the conjectured solution back into equation (35) one obtains the stated expression for ω_h. ■

From the Growth Option Lemma we obtain the following expressions for the growth option value function during the Implementation Stage:

$$
\begin{aligned}
G^I(x) &= \sum_{h=0}^{\nu} \phi_h^I \omega_h^I x^h \qquad (36)\\
\text{with } \phi_h^I &\equiv \binom{\nu}{h} (\eta^I)^{\nu-h} \zeta^h \Gamma.\\
\omega_h^I &\equiv \frac{1}{r - \mu h - \frac{1}{2}\sigma^2 h(h-1)}.
\end{aligned}
$$

Experiment Stage

The HJB equation for the Experiment Stage is:

$$
\begin{aligned}
rV^E(x,k) &= \max_i \quad (x+\pi_E b)k - \gamma i^{\nu/(\nu-1)} + \mu x V_x^E(x,k) + \frac{1}{2}\sigma^2 x^2 V_{xx}^E(x,k) \qquad (37)\\
&\quad + (i-\delta k)V_k^E(x,k) + \lambda_I \beta \left[\overline{V}^I(x,k) - V^E(x,k)\right] + \lambda_I(1-\beta)\left[\underline{V}^I(x,k) - V^E(x,k)\right].
\end{aligned}
$$

We conjecture and verify the value function is separable between the value of assets in place and growth options:

$$
V^E(x,k) = kq^E(x) + G^E(x). \qquad (38)
$$

Under the posited functional form, the optimal instantaneous control during the Experiment Stage is $i^*(q^E)$. Substituting the Implementation Stage value functions into the HJB equation, and isolating the terms scaled by k, we obtain the following ODE for the shadow value of capital:

$$
(r+\delta+\lambda_I)q^E(x) = (x+\pi_E b) + \mu x q_x^E(x) + \frac{1}{2}\sigma^2 x^2 q_{xx}^E(x) + \lambda_I \left[\frac{x}{r+\delta-\mu} + \frac{(\beta\overline{\pi} + (1-\beta)\underline{\pi})b}{r+\delta}\right]. \qquad (39)
$$

We conjecture the following linear form for the shadow value of capital:

$$
q^E(x) = x\zeta^E + \eta^E. \qquad (40)
$$

Substituting the conjectured solution into the ODE for q, we obtain the following solution, as presented in the body of the paper:

$$
\begin{aligned}
\zeta^E &= \zeta \equiv 1/(r+\delta-\mu). \qquad (41)\\
\eta^E &= \left[\frac{(r+\delta)\,\pi_E b + \lambda_I(\beta\overline{\pi} + (1-\beta)\underline{\pi})b}{(r+\delta)\,(r+\delta+\lambda_I)}\right].
\end{aligned}
$$

We next determine the growth option value for the Experiment Stage. Proceeding as above and dropping the terms scaled by k in the HJB equation, we obtain the condition:

$$
(r+\lambda_I)\,G^E(x) = \mu x G_x^E(x) + \frac{1}{2}\sigma^2 x^2 G_{xx}^E(x) + \lambda_I\left[\beta\overline{G}^I(x) + (1-\beta)\underline{G}^I(x)\right] + \sum_{h=0}^{\nu}\left[\binom{\nu}{h}\left(\eta^E\right)^{\nu-h}\zeta^h\Gamma\right]x^h. \qquad (42)
$$

This ODE can be rewritten as:

$$(r+\lambda_I)\,G^E(x) = \mu x G_x^E(x) + \frac{1}{2}\sigma^2 x^2 G_{xx}^E(x) + \sum_{h=0}^{\nu} \phi_h^E x^h. \tag{43}$$

where

$$\begin{aligned}
\phi_h^E &\equiv \left[\lambda_I\left(\beta\overline{\phi}_h^I + (1-\beta)\underline{\phi}_h^I\right)\omega_h^I\right] + \binom{\nu}{h}\left(\eta^E\right)^{\nu-h}\zeta^h\Gamma \\
\overline{\phi}_h^I &\equiv \binom{\nu}{h}(\overline{\eta}^I)^{\nu-h}\zeta^h\Gamma \\
\underline{\phi}_h^I &\equiv \binom{\nu}{h}(\underline{\eta}^I)^{\nu-h}\zeta^h\Gamma.
\end{aligned}$$

From the Growth Option Lemma it follows:

$$\begin{aligned}
G^E(x) &= \sum_{h=0}^{\nu}\phi_h^E\omega_h^E x^h \\
\text{with } \omega_h^E &\equiv \frac{1}{(r+\lambda_I)-\mu h-\frac{1}{2}\sigma^2 h(h-1)}.
\end{aligned}$$

Pre-Experiment Stage

The HJB equation is:

$$\begin{aligned}
rV^P(x,k) &= \max_i \quad (x+\pi_P b)k - \gamma i^{\nu/(\nu-1)} + \mu x V_x^P(x,k) + \frac{1}{2}\sigma^2 x^2 V_{xx}^P(x,k) \\
&\quad +(i-\delta k)V_k^P(x,k) + \lambda_E\theta\left[\overline{V}^E(x,k) - V^P(x,k)\right] + \lambda_E(1-\theta)\left[\underline{V}^E(x,k) - V^P(x,k)\right].
\end{aligned} \tag{44}$$

We again conjecture a value function separable between the value of assets in place and growth options:

$$V^P(x,k) = kq^P(x) + G^P(x). \tag{45}$$

Inspecting the HJB equation it is apparent that the optimal control policy during Stage P is $i^*(q^P)$.

Substituting the conjectured value function into the HJB equation and isolating those terms scaled by k, we obtain the following ODE for the shadow value of capital:

$$(r+\delta+\lambda_E)q^P(x) = (x+\pi_P b) + \mu x q_x^P(x) + \frac{1}{2}\sigma^2 x^2 q_{xx}^P(x) + \lambda_E\theta\overline{q}^E(x) + \lambda_E(1-\theta)\underline{q}^E(x). \tag{46}$$

We may again conjecture (and verify) the preceding shadow value equation has a linear solution, resulting in equation (20).

We turn next to determining growth option value during the Pre-Experiment Stage. Confining attention to the remaining terms in the HJB equation that are not scaled by k, we obtain the

following ODE:

$$\begin{aligned}(r+\lambda_E)G^P(x) &= \mu x G_x^P(x) + \frac{1}{2}\sigma^2 x^2 G_{xx}^P(x) + \sum_{h=0}^{\nu}\left[\binom{\nu}{h}(\eta^P)^{\nu-h}\zeta^h\Gamma\right]x^h \\ &\quad +\lambda_E\left[\theta\overline{G}^E(x) + (1-\theta)\underline{G}^E(x)\right].\end{aligned} \tag{47}$$

Substituting in the expressions for G^E and grouping terms one obtains:

$$(r+\lambda_E)G^P(x) = \mu x G_x^P(x) + \frac{1}{2}\sigma^2 x^2 G_{xx}^P(x) + \sum_{h=0}^{\nu}\phi_h^P x^h. \tag{48}$$

with

$$\begin{aligned}\phi_h^P &\equiv \binom{\nu}{h}(\eta^P)^{\nu-h}\zeta^h\Gamma + \lambda_E\omega_h^E\left[\theta\overline{\phi}_h^E + (1-\theta)\underline{\phi}_h^E\right] \\ \overline{\phi}_h^E &\equiv \left[\lambda_I\left(\beta\overline{\phi}_h^I + (1-\beta)\underline{\phi}_h^I\right)\omega_h^I\right] + \binom{\nu}{h}\left(\overline{\eta}^E\right)^{\nu-h}\zeta^h\Gamma \\ \underline{\phi}_h^E &\equiv \left[\lambda_I\left(\beta\overline{\phi}_h^I + (1-\beta)\underline{\phi}_h^I\right)\omega_h^I\right] + \binom{\nu}{h}\left(\underline{\eta}^E\right)^{\nu-h}\zeta^h\Gamma.\end{aligned} \tag{49}$$

Again, the growth option value is a linear sum of the geometric Brownian motion x to successive powers. From the Growth Option Lemma it follows that growth option value during the Pre-Experiment Stage is:

$$\begin{aligned}G^P &= \sum_{h=0}^{\nu}\phi_h^P\omega_h^P x^h \\ \omega_h^P &\equiv \frac{1}{r+\lambda_E-\mu h-\frac{1}{2}\sigma^2 h(h-1)}.\end{aligned} \tag{50}$$

Figure 1: Endogenous Beliefs

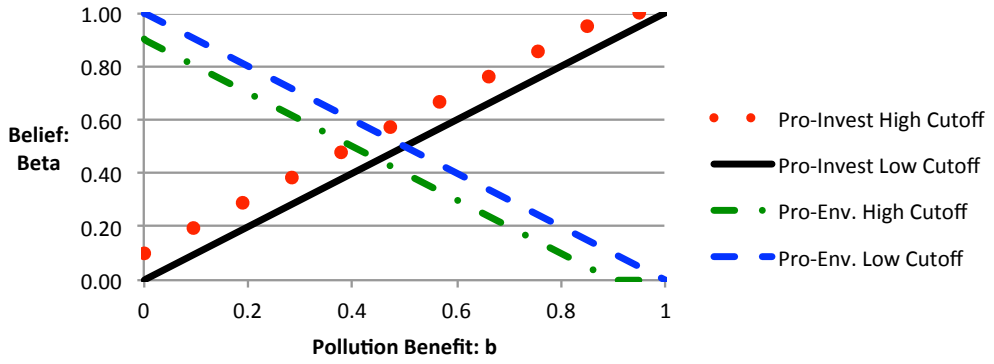

Figure 2: Natural Policy Experiment

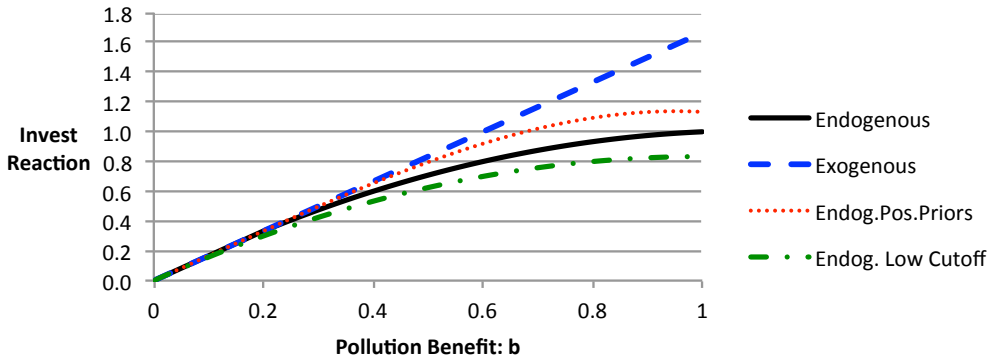

Figure 3: Hawthorne Effects in NPEs

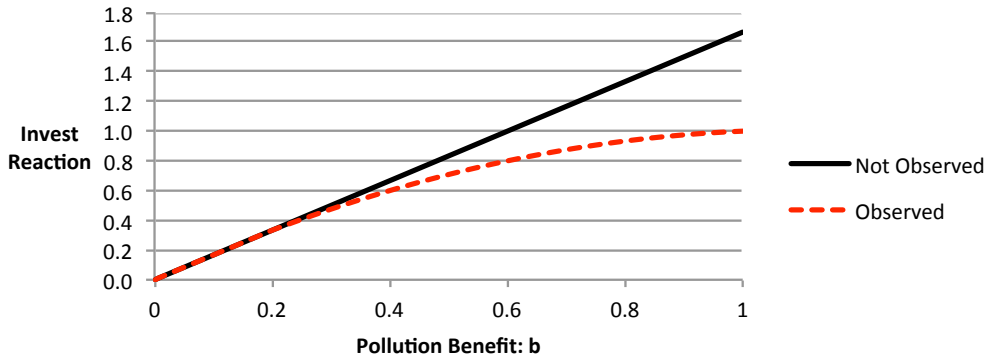

Figure 4: Randomized Controlled Trial

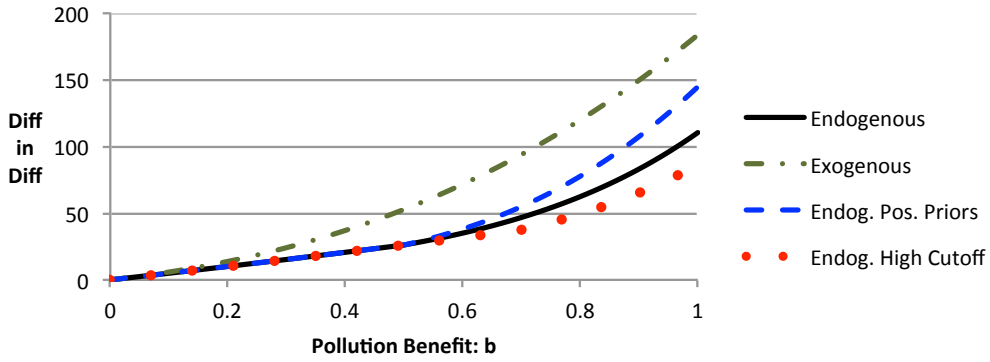

Figure 5: Observer Effects in RCTs

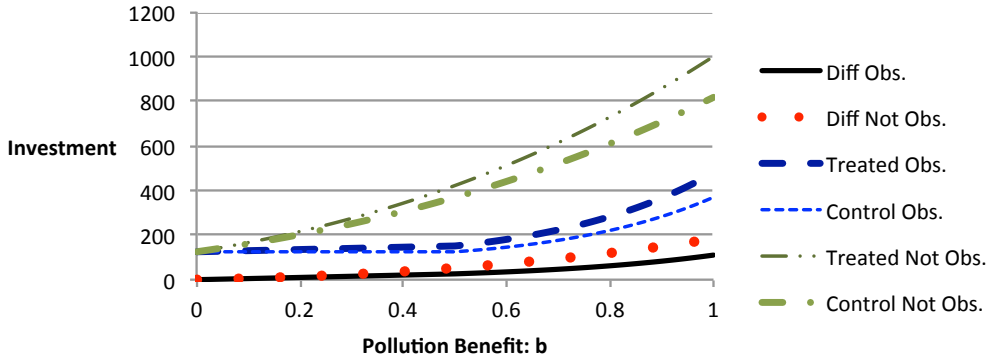

Figure 6: Irreversibility and RCTs

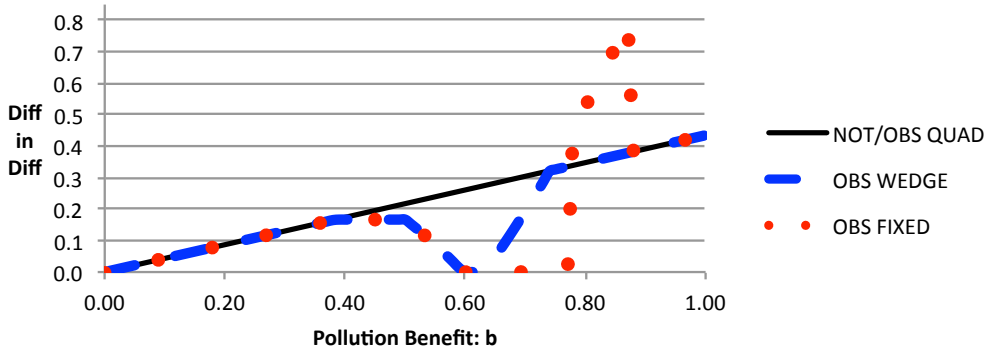

Figure 7A: NPE Identification Failure

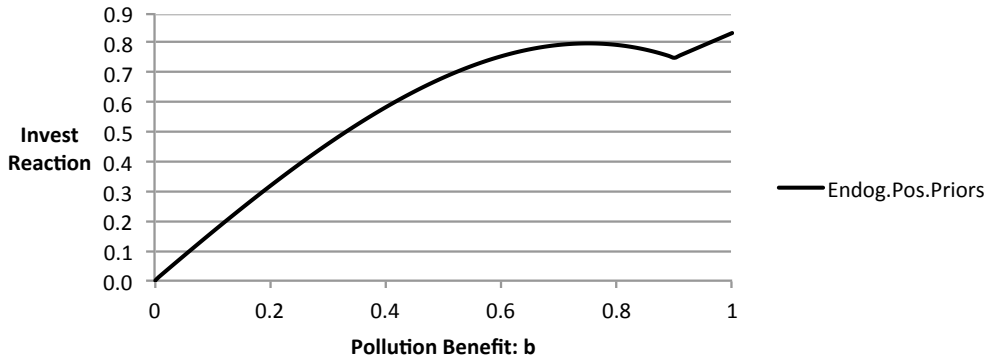
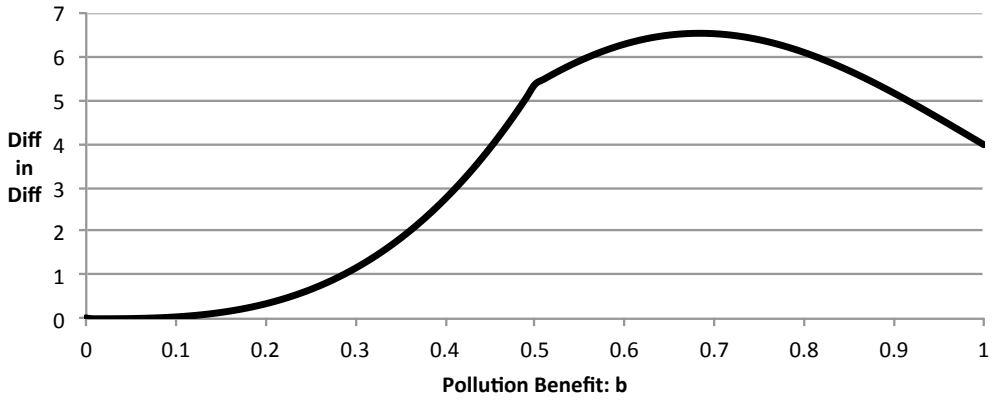
Figure 7B: RCT Identification Failure