The Paradox of Policy-Relevant RCTs and Natural Experiments*

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

Christopher A.Hennessy
LBS, CEPR, and ECGI

October 2016

Abstract

According to conventional wisdom, natural experiments represent especially credible bases for causal parameter inference, facilitating evidence-based policymaking. We examine robustness of evidence derived from first-stage randomizations applied to measure zero subjects. If government can (cannot) alter policy in response, experimental evidence is contaminated (uncontaminated) by *ex post policy endogeneity*: Measured responses depend upon prior beliefs and the government objective function into which evidence is fed. If government perceives experimental evidence as credible (non-credible), 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. Endogenous belief heterogeneity arises from heterogeneous causal effect parameters, potentially preventing parameter inference. Treatment-control differences in RCTs are contaminated unless stock variable accumulation cost functions satisfy: 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

In their influential textbook, *Mostly Harmless Econometrics*, Angrist and Pischke (2009) argue empirical evidence delivered via quasi-randomization represents a credible stand-alone product: "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." Echoing this view, Greenstone (2009) writes, "The gold standard for estimating the causal impact of a regulation is the randomized trial." Bowen, Frésard, and Taillard (2016) note that, "The use of such techniques has recently become a widespread tool for corporate finance researchers, mirroring the trend observed in other areas of economics (e.g. labor and development)." Some micro-econometricians have claimed the scientific high ground, with Angrist and Pischke (2010) and Romer (2016), amongst many others, criticizing macroeconomists and asset pricers for allegedly putting calibration before causation.

Ingenuity notwithstanding, the "credibility revolution" heralded by Angrist and Pischke (2010) has faced criticism for allegedly 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 created a trend toward its direct use in setting policy. For example, as discussed by Spatt (2011), the SEC has begun to use evidence from random assignment to inform its decisions, having met judicial challenges to its cost-benefit analyses. More generally, Dhaliwal and Tulloch (2015) note the existence of an "increasing trend towards considering rigorous evidence while making policy decisions." Greenstone (2009) calls for government "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."

Clearly, a correct assessment of what quasi-natural experiments can and cannot achieve is of great importance in light of the strong methodological and epistemological claims made by some in the randomization camp, and the apparent traction these claims hold with policymakers and young researchers. To this end, this paper describes a logical inconsistency at the heart of the natural experiment research program as it relates to inference in *dynamic* settings– the very types of settings likely to be of interest in finance, where the object of empirical study and governmental regulation is often an infinitely-lived corporation. A simple example illustrates the root problem. We begin by noting that the two primary, and often exclusive, objectives of contemporary empirical work are to convince the audience of clean identification (allegedly achieved via randomization) and policy relevance. Suppose then, that after exhaustive debate, which is the current norm in empirical due diligence, an empiricist is able to convince his audience that it was Nature herself that forced an exogenous change in government policy or that Nature randomly generated treatment and control groups. The empiricist is next challenged on policy-relevance. Suppose that here too he can rise to the challenge, e.g. the government has rationally decided to use "credible estimates" of causal parameters as inputs into future regulatory cost-benefit analyses. 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 contradiction between the 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, the empiricist has actually 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, they will have rationally changed their behavior during the experiment in light of the anticipated influence of econometric analysis. That is, rational anticipation of evidence-based policymaking post-experiment changes what the econometrician measures during the experiment. What implications does this have for causal inference?

To examine this question formally, we consider the following setting. At each point in time, agents (firms) make optimal (investment) decisions in light of current and expected future regulation, modeled as a cap on pollution. Econometricians and the government would like to infer the marginal benefit each industry derives from relaxation of pollution regulation, since this determines the key causal effect parameter in our economy: the effect of pollution caps on long-term investment. In the model, marginal pollution benefits are private information, modeled as industry-specific i.i.d. draws from a known distribution.

Fortunately for econometricians in our economy, 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 are 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 face regulation during the Experiment Stage and the remaining firms do not. In the NPE setting, econometricians attempt to infer industry-level pollution benefits based upon investment responses to the regulation shock. In the RCT setting, inference is instead 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, considering three post-experiment scenarios. In the first scenario, the government is powerless to change the policy variable post-experiment. In the second scenario, the government is able to change the policy variable post-experiment, and will do so using the experimental evidence. For example, a pro-investment government might deregulate only if inferred 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.

As we show, 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 a seemingly ideal first-stage randomization is contaminated by what we term *ex post endogeneity.* And this is true even if, as in our economy, individual firms are measure zero and have no ability to influence empirical tests or policy decisions. Importantly, the problem of ex post endogeneity vanishes if the government is powerless to change future policy. The following paradoxical situation thus emerges: The experimental evidence is uncontaminated only if the government is unable to use it. Similarly, contamination from ex post endogeneity vanishes if the government does not view the experimental evidence as credible and ignores it. Thus, another paradoxical situation emerges: The experimental evidence is uncontaminated only if the government does not deem it to be credible.

We move well beyond illustrating these paradoxes, describing five novel challenges to causal parameter inference arising from ex post endogeneity. First, rather than being a stand-alone object, 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 first-stage policy randomization which, by definition, does not depend on government objective function parameters, does not eliminate the need to make assumptions about these same parameters. After all, the government's objective function will influence the distribution of the policy variable after the experiment, thus influencing measured responses during the experiment. In other words, policy endogeneity (selection), and its ability to cloud causal inference, has only been pushed back in time.

Second, with policy-relevant experimentation, causal parameters can only be correctly inferred if one has correctly stipulated the prior beliefs held by agents regarding the probability distribution of these same causal parameters. Intuitively, since agents know government will base its post-experiment policy on causal parameter estimates, prior beliefs regarding the probability distribution of these causal parameters influence agent beliefs regarding the distribution of the policy variable post-experiment. This influences the measured responses of forward-looking agents during the experiment. It follows that an incorrect stipulation of prior beliefs regarding causal parameters leads to incorrect inference of these same parameters.

Third, the ex post endogeneity problem described above generates observer effects: the act of observation by econometricians changes the measured responses of both treatment and control groups. And this is true despite our agents being rational, anonymous, and measure zero. Intuitively, variation in the post-experiment policy variable distribution, resulting from econometric observation cum evidence-based policy-setting, changes measured responses during the experiment.

Fourth, the observer effects described in the preceding paragraph are unequal across treatment and control groups unless the underlying stock variable accumulation technology satisfies the type of "strong functional form assumptions" that randomization advocates have criticized in structural econometric estimations (e.g. Hayashi (1982)): zero fixed costs, equality of buy and sell prices, and quadratic adjustment costs. Consequently, if these functional form assumptions are not satisfied, treatment-control differences in RCTs are contaminated by observer and policy feedback effects, with incorrect causal parameter inference resulting if they are not taken into account.

Fifth, endowed heterogeneous causal effect parameters across agents generate endogenously heterogeneous beliefs regarding post-experiment policy. That is, if there is cross-sectional variation in causal effect parameters, a common assumption in applied micro-econometric work, then there will be concomitant cross-sectional heterogeneity in policy expectations. Since, as we show, this beliefs channel either amplifies or attenuates treatment response heterogeneity, a failure to take it into account leads to faulty inference regarding the magnitude of causal parameters.

Finally, the fact that beliefs are a function of causal parameters can make it impossible to recover causal parameters from RCT and NPE test statistics. Beliefs and causal parameters become confounded. Formally, the potential impossibility of recovering causal parameters from test statistics is due to the fact that endogenous policy beliefs can cause experimental moments, e.g. the treatment-control difference, to be non-monotone in causal effect parameters, so that the moments cannot be inverted to solve for the true value of the causal parameter. The intuition is as follows. Consider a measure zero firm whose industry stands to benefit a great deal if higher pollution is permitted. Such a firm would tend to increase investment greatly if there were a pollution deregulation experiment. However, this same firm might actually not respond much at all during the experiment if it rationally conjectures that the government is more likely to regulate post-experiment after learning each industry's causal parameter (including those of its own high-pollution industry).

In the model, the agents exposed to the first-stage randomization make decisions during the experiment (as opposed to crops, cells, or particles), and these 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, the model applies to a broad range of empirical settings involving the accumulation of stock variables: financial structure, employees, savings, debt, health, human capital, offspring, and reputation in principal-agent settings. Further, the model is applicable to many superficially static decisions once one accounts for life-cycle considerations, e.g. labor supply. In the model, the agents being measured face endogenous government policy responses to the experimental evidence. We close the paper (Section 6) by providing numerous real-world examples of such policy feedback. 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 nexus between estimation and policy-setting, a nexus that many randomization advocates promote.

The issues raised are related to, but distinct from, the econometric critique made by Lucas (1976). Writing for *New Palgrave Dictionary of Economics*, Ljungqvist (2008) defines the Lucas Critique as follows:

> 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 how experimental evidence 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 post-experiment policy change that is the root cause of the five novel econometric challenges we demonstrate, especially the role of the government objective function, the role of prior beliefs regarding the causal parameters to be estimated, and the confounding effects of endogenous policy belief heterogeneity. 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, their perceived credibility, and government policy. This feedback is an underlying cause of the novel biases and paradoxes we illustrate, upon which Lucas (1976) is silent. These differences notwithstanding, the present paper borrows from Lucas the idea of viewing empirical tests, here policy-relevant experiments, through the prism of rational expectations.

Bond, Goldstein and Prescott (2009) consider the related but distinct issue of how the use of securities prices in setting policy can change price informativeness. For example, incentives to acquire information can be attenuated if government may take actions rendering the information irrelevant for securities payoffs. Our argument does not concern informed trading or securities prices, but rather centers on the correct interpretation of empirical test statistics, e.g. measured investment changes, derived from seemingly-ideal first-stage policy randomizations. In their model, feedback from securities prices to policy limits price informativeness. In our model, information quality need not suffer provided the econometrician correctly accounts for the policy feedback loop.

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. 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 governments. In contrast, we analyze the correct interpretation of experimental evidence assuming all agents, including the government, behave optimally and make optimal use of their information.

Our 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 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 paradoxes 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 violate random assignment. In our laboratory economy, firms are incapable of avoiding the experimental treatments. Heckman (1997) and Deaton (2010) emphasize that with heterogeneous causal effect parameters, 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 problems associated with small samples and bias in panel selection. We consider infinite sample sizes, in that there is a continuum of treated and control firms, with ideal first-stage policy randomization.

Acemoglu (2010) argues general equilibrium effects can limit the external validity of small-scale experiments. In particular, he argues large-scale policy changes potentially lead to factor substitution and endogenous changes in prices and technologies. These effects are shut off in our model. Acemoglu also points out that endowed 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.

Chassang, Padro i Miguel and Snowberg (2012) consider static RCTs and show how hidden effort during an experiment can cloud inference regarding treatment efficacy. For example, low average treatment effects can arise from truly low efficacy or low agent effort caused by erroneous expectations of low treatment efficacy. Our model abstracts from hidden effort and their model abstracts from endogenous post-experiment policy, so the bias causes differ fundamentally. In their model, beliefs concern treatment efficacy, not the stochastic path of long-term policy variables. Thus, the essential point and paradox in our paper, that evidence-based long-term policymaking implies violation of the standard treatment orthogonality assumption, and clouded causal parameter inference, is necessarily absent from their paper.

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. The behavioral and organization literatures have postulated a range behavioral rationales for such effects such as self-consciousness, approval-seeking, spite, or a desire to influence study outcomes. Importantly, our model abstracts from each of these effects since firms are rational, measure zero, and anonymous.

The rest of the paper is as follows. Section 2 presents a model of the interaction between firms, governments, and econometricians. Section 3 discusses econometric inference in settings where firms face a common exogenous policy shock (NPEs). Section 4 discusses RCTs. Section 5 illustrates how the inference-policy loop can cause outcome variables to become non-monotone in causal parameters, blocking causal parameter inference. Section 6 provides real-world examples of the experimental inference-policy loop.

2 The Model

We begin by contrasting inference in two economies endowed with identical natural experiments and technologies but differing in whether the empirical evidence will be used. In the *Endogenous Policy Economy*, the experimental evidence will be used to select an optimal policy post-experiment. In the *Exogenous Policy Economy*, the experimental evidence is irrelevant because the government is powerless to change the path of the policy variable. The investment model is basic, following, say, Dixit and Pindyck (1994).

2.1 Technology

Time is continuous and the horizon infinite. Agents are risk-neutral and share a common discount rate $r > 0$. There is a measure one continuum of anonymous firms with generic member $j \in \mathbb{J}$. Since firms are atomistic, no firm has any incentive to change its behavior with the goal of influencing test statistics, econometric inference, or government policy. That is, each firm acts as a policy-taker.

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 firm invests optimally each instant.

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 the main arguments, we start by assuming firms face the following convex investment cost function:

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

We assume the cost function parameters satisfy $\gamma > 0$ and $\nu > 1$. Imposing $\nu > 1$ ensures unique optimal policies. Further, here one obtains simple closed-form expressions for the empirical outcome variable i. Since the firm's value function is of peripheral interest, we relegate its derivation to the appendix, confining attention to integer values of ν, as in Abel and Eberly (1997).

Firm cash-flow cannot be observed by the government. Cash-flow at date 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, assume the discount rate satisfies $r > \mu\nu + \frac{1}{2}\sigma^2\nu(\nu - 1)$.

We shall think of all firms in the economy as facing the same x process. This assumption is not necessary, but serves the expositional purpose of approximating the type of macroeconomic shocks to which a real-world government might be expected to respond. Analogous sources of endogeneity bias are generally of paramount concern to econometricians. For example, empiricists are often concerned about downward bias in causal parameter estimates if governments are more willing to impose regulations in good times (high x). Anticipating, the exogenous first-stage policy randomizations we consider are such that this and other standard forms of (ex ante) endogeneity and selection bias will not be an issue– the experiments considered will be optically ideal.

The term $\pi_t b$ entering cash-flow captures government policy. The variable π represents a government-stipulated limitation on some privately profitable activity generating a negative social externality. Externality generating activities are ubiquitous in finance: welfare-reducing contract designs (e.g. Hart (2009)); unstable bank asset and liability structures (e.g. Gorton (2010)); quote-stuffing by high-frequency traders (e.g. Biais and Woolley (2011)); or compensation-increasing governance structures (e.g. Acharya and Volpin (2010)). To fix ideas, we call this regulated activity "pollution."

The variable $b \geq 0$ measures an unobservable private benefit to polluting, e.g. cost or effort reductions. This benefit can be pecuniary or non-pecuniary. Pollution benefits are industry-specific, and there are $M \geq 2$ industries. Econometricians know each firm's industry, but do not know the marginal benefit, call it b_m, firms in industry m derive from a higher pollution cap. The key ingredient here is that firm technologies are correlated. Therefore, a firm's own technology is informative about what econometricians will measure during the experiment, thus helping to predict endogenous government policy decisions post-experiment.

The industry-specific pollution benefits are i.i.d. draws at date 0. The pollution benefit for an industry is drawn from the interval $[0, \bar{b}]$ 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 want to infer the realized vector of industry-level pollution benefit parameters $\mathbf{b}$. For example, in the Endogenous Policy Economy, this inference will inform decisions regarding the optimal pollution cap post-experiment. We thus speak of F as capturing prior beliefs over the causal parameters.

2.2 Timing

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 $\overline{\pi}$, with the binary policy assumption simplifying the econometric inference problem. 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. During each stage, firms face a simple time-homogeneous investment problem, so we essentially have a three period 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 $\widetilde{\tau}_E$. 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 $\pi_E = \overline{\pi}$, with the remaining firms 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 an NPE, the model parameter θ is set to 0 or 1. In a *Randomized Controlled Trial* (RCT below), treatment and control groups face different regulations during the experiment. To capture an RCT, the parameter θ is chosen from the interval $(0, 1)$. For example, in a standard RCT featuring an equal measure of firms in treatment and control groups, $\theta = 1/2$.

"Unexpected" experiments are captured 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 regarding biased inference.

The Implementation Stage I will arrive at date $\widetilde{\tau}_I$. This date is an independent random variable given $\widetilde{\tau}_E$. The transition rate into the Implementation Stage is $\lambda_I > 0$. Thus, at any instant prior to such a 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 an RCT, econometricians can look back and measure the difference between the investment of treatment and control groups, industry-by-industry, at some date during the experiment. In an NPE, econometricians can look back and measure the jump in each industry's investment that occurred at the start of the experiment. Since the path of the macroeconomic profit factor x is continuous, measuring investment changes over the instants just before and just after the experiment is initiated eliminates the need to control for changes in macroeconomic conditions (the state variable x here). As shown below, the parameter vector of interest $\mathbf{b}$, 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): $\widetilde{\pi}_{jt} \perp \{\widetilde{\mathbf{b}}, \widetilde{x}_t\} \quad \forall \quad j \in \mathbb{J}$ and $t \in [0, \widetilde{\tau}_I)$.

By construction, independence of the policy process rules out the standard forms of endogeneity bias about which empiricists would be expected to debate. 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 experiment with the novel policy (π_E) because it knows that the true causal parameter configuration in its economy is favorable to the policy. However, in the economy considered, the government does not know the parameter

[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.

vector $\mathbf{b}$, and the policy variable is independent of the random vector $\widetilde{\mathbf{b}}$ at any point in time at which the experimental measurement may take place, specifically all times $t \in [0, \widetilde{\tau}_I)$. Finally, recall from equation (4) that the Wiener process w_t was assumed to be independent, eliminating concern that experimental regulation is correlated with the underlying profitability state (x_t). In this way, the experiment will have ideal optics.

Since we will be evaluating the difference-in-difference estimator, it is worth noting that here too our experiments will have ideal 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) prior to the experiment.[3] In particular, for all firms j within a given industry m we have:

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

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

During the Implementation Stage, a long-term regulatory policy (π_I) will be implemented permanently. In the Exogenous Policy Economy, π_I will be set to a technologically pre-determined 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 implement an optimal regulatory policy in light of this information.

2.3 Endogenously Heterogeneous Policy Beliefs

When the Implementation Stage is reached in the Endogenous Policy Economy, the government will put into place an optimal policy according to an objective function Θ. The objective function is common knowledge to all agents at date 0. For simplicity, assume the government will choose policy based on average industry-level pollution benefits. The government's "type" is fixed, and it can be either Pro-Environment or Pro-Investment, with the type being common knowledge. The corresponding objective functions are as follows.

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

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

In the preceding equations, $b^* > 0$ is the key parameter in the governmental objective function determining policy decisions. Specifically, a Pro-Investment government will deregulate if and only

[3] In fact, the trend is common across all firms and industries in this model.

if the average industry-level pollution benefit exceeds b^*. Intuitively, the long-term (Implementation Stage) investment of each firm will be increasing in its pollution benefit parameter b_j, and a Pro-Investment government will deregulate if doing so will bring about a sufficient investment stimulus. If the government is instead Pro-Environment, it will regulate if and only if the average industry-level pollution benefit exceeds b^*. Intuitively, if the average pollution benefit is high, deregulation would lead to large investments in polluting capital stocks, investments a Pro-Environment government wants to deter.

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

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

The following parametric assumption ensures that econometric inference actually serves a real purpose in the Endogenous Policy Economy.

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

It is readily verified that if the preceding condition were not satisfied, then the government would never change its policy regardless of the econometric evidence.

Let χ be an indicator function equal to 1 if deregulation is optimal:

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

The policy beliefs for a firm in industry m are captured by a function β capturing their assessment of the probability of deregulation during the Implementation Stage. We have:

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

Notice, under the stated assumptions, the belief function β is the same across industries. However, the realized value of its argument, the own-industry pollution benefit b, is different across industries with probability one.

From equation (8) it follows that the endogenous government policy decision during the Implementation Stage is necessarily uncertain, with the inferred causal parameter value for each industry influencing marginal decisions. Further, in the case of a Pro-Investment (Pro-Environment) government, it is apparent that the indicator function in equation (10) is non-decreasing (non-increasing) in a firm's respective own-industry pollution benefit parameter b. This implies that under a Pro-Investment (Pro-Environment) government, a firm experiencing a high b value will think it more likely for the government to deregulate long-term. In this way, heterogeneous causal effect parameters lead to endogenously heterogeneous policy beliefs. The following lemma summarizes.

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

It is useful to close this subsection by describing beliefs in the simplest possible setting where the economy has only two industries ($M = 2$). From equation (10) it follows that with two industries, beliefs take the following form.

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

It is apparent from the preceding equations that, consistent with the preceding lemma, endowed heterogeneity in treatment response parameters (b_1, b_2) results in endogenously heterogeneous policy beliefs $(\beta(b))$ across industries. Further, the preceding equation reveals that the shape of the belief function is determined by prior beliefs regarding the parameters to be estimated (F), as well as the government objective function parameter b^*.

Figure 1 plots realized beliefs as a function of the realized own-industry pollution benefit parameter b. The figure assumes F is the uniform distribution on $[0, 1]$ and plots beliefs under both Pro-Investment and Pro-Environment governments. Consistent with the preceding lemma, the belief function is increasing in b under the former and decreasing in b under the latter. The figure considers governments applying both 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).

A number of points are worth stressing before closing this subsection. First, it may be tempting to argue that correct inference boils down to having a correct model of the policy variable distribution, e.g. making a correct assumption of the probability of policy permanence. However, the preceding lemma shows that there is no such thing as a common policy expectation. Cross-sectional differences in causal parameters generate endogenously heterogeneous policy beliefs if the government engages in evidence-based policymaking. Second, assuming that one can make correct *a priori* assumptions about each industry's policy belief $(\beta(b))$ is tantamount to assuming one knows the causal parameters $(\mathbf{b})$ of interest. But if this were indeed the case, the whole econometric exercise would be unnecessary.

Although the details will differ, the type of endogenously heterogeneous policy expectations derived above are likely 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 assumes that the government, capable of correct statistical inference, is more likely to continue or scale-up (discontinue or scale-down) the policy. As shown below, a failure to account for this beliefs channel will lead to incorrect inference regarding causal parameters in NPEs (Section 3) and RCTs (Section 5). Worse still, the fact that beliefs are endogenous functions of causal parameters may actually prevent the recovery of causal parameters (Section 7).

2.4 Investment Decisions

The model is formally solved in the appendix via optimal control. 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 is thus solved as readily as a three period model.

We begin by focusing on the Endogenous Policy Economy. At each instant, it is optimal for the firm 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. Therefore, the following function i^* maps the instantaneous shadow value of capital 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.

Since 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 1 unit of capital which, in turn, is just 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 of the flow of future pollution benefits in state S for a firm with pollution benefit b. This pollution benefit value function is derived next via 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 thus 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 present 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. Thus, we have the following equilibrium condition:

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

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 present 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:

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

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

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

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

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

Effectively, in the Exogenous Policy Economy, beliefs are homogeneous across firms and industries, 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:

$$CE_m \equiv \frac{\partial i^I(x, b_m; \pi)}{\partial \pi} = \frac{\partial q^I(x, b_m; \pi)}{\partial \pi}\frac{di^*}{dq} = \left(\frac{b_m}{r+\delta}\right)\frac{di^*}{dq}. \tag{22}$$

Notice, the preceding equation implies that correctly forecasting firm responses to the final governmental policy decision (π_I), requires correctly inferring the vector $\mathbf{b}$ of causal effect parameters. The rest of the paper considers this inference problem.

3 Inference in Natural Policy Experiments

This section considers causal parameter inference in the context of Natural Policy Experiments (NPEs) in which all firms face a common exogenous policy change. Recall, from equation (22), we know the causal parameters of interest here are the vector $\mathbf{b}$ of industry-specific pollution benefits which determine the investment response to long-term (Implementation Stage) deregulation.

3.1 Barriers to Inference in NPEs

The econometric challenges in NPEs are perhaps best understood by considering a specific NPE. To this end, suppose that during the Pre-Experiment Stage pollution is regulated, with $\pi_P = \underline{\pi}$. During the Experiment Stage all firms will be deregulated, with $\pi_E = \overline{\pi}$. In the Exogenous Policy

Economy, the policy variable will remain fixed ($\pi_I^{EX} = \overline{\pi}$) after the experiment and the government is powerless to change it. In the Endogenous Policy Economy, the government's econometrician will correctly infer the vector $\mathbf{b}$ based on the industry-level responses to the policy experiment. The government then implements its optimal policy. In this setting, we evaluate inference by an econometrician sitting outside the government, say, an academic economist.

The numerical examples in this section consider that there are two industries, with their pollution benefits (b) being i.i.d. draws from the uniform distribution on $[0, 1]$. In the Endogenous Policy Economy, there is a Pro-Environment government in charge. This government will deregulate only if the average pollution benefit falls below the cutoff value $b^* = 0.60$. 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 $\nu = 2$. Notice, we here consider long expected policy regime durations of $\lambda^{-1} = 6.7$ years, approximating reality for economy-wide policy legislation. It is also worth noting that the numerical examples also assume adjustment costs are quadratic ($\nu = 2$), implying the optimal investment mapping (equation (12)) is linear in q.

The issues here are illustrated in Figure 2. On the horizontal axis in the figure is the true pollution benefit for one of the two industries, the parameter our econometrician wants to infer. On the vertical axis is the corresponding increase in investment that will occur at the start of the experimental deregulation period, which depends upon the true magnitude of the causal parameter. Consider first the contrast between the responses in the Endogenous versus Exogenous policy economies. In both cases, there is zero investment response if $b = 0$. After all, if cash-flow is invariant to the pollution cap, investment will also be invariant to the cap. Further, in the present numerical example, the investment response is monotonically increasing in b. Monotonicity implies that our econometrician will be able to correctly infer the causal parameter's value provided that she correctly accounts for the interplay between estimation, expectations, and policy-setting.

Suppose now our academic econometrician works in the Endogenous Policy Economy– the type of economy envisioned by randomization advocates, one in which policy is "evidence-based." 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 feedback role is ignored and she instead uses the dashed line to perform inference. Notice, the dashed line represents a counter-factual economy in which the experimental policy change is permanent with probability 1, so that causal parameter inference is policy-irrelevant. From Figure 2 it is apparent the naïve econometrician will understate the causal parameter. For example, if $b = 1$ the observed investment increase will be equal to 1. However, working along the dashed line, the naïve econometrician will infer $b = .60$. Intuitively, in the Endogenous Policy Economy, the government can 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 for any given value of b.

Figure 2 also reveals that the wedge resulting from the policy expectations effect varies with the true value of the causal parameter. In particular, the gap between the solid and dashed schedules grows larger as b grows larger. As shown formally below, this non-linearity is due to the fact that a higher own-industry pollution benefit causes a firm to rationally infer a lower probability of deregulation long-term, consistent with our demonstration that under a Pro-Environment government $\beta' \leq 0$. In other words, heterogeneous causal parameters lead to heterogeneous policy beliefs, affecting what is measured. Here the failure to account for the compression of treatment responses across b values, resulting from endogenous belief heterogeneity, apparently leads to downward bias.

Another issue of concern to econometricians is that selection limits external validity. For example, if a Pro-Environment 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 experimental deregulation considered here. But recall, throughout the paper we are adopting Assumption 1 which rules out this standard form of selection. However, ex post endogeneity will lead to similar limits to external validity. For example, if two economies simply differ in terms of the prior beliefs held by agents regarding the probability distribution of causal parameters (F), their treatment response functions will differ. It follows that correct causal parameter inference requires correctly specifying prior beliefs over the causal parameter distribution.

To illustrate, suppose that in addition to the baseline Endogenous Policy Economy (solid line), there is an otherwise equivalent economy in which agents hold more favorable prior beliefs (from the perspective of a Pro-Environment government), with F now being a triangular distribution placing more weight on lower b values. This new case is captured by the dotted line in Figure 2. Notice, for any given value of the causal parameter b, the investment reaction at inception of the experiment will actually be stronger in the economy with more favorable priors. Intuitively, with more favorable priors regarding the causal parameter to be estimated, firms assign higher probability to deregulation post-experiment, and so invest more. The necessity of common priors is a severe limit on external validity, and the need to correctly specify priors is a severe challenge to causal parameter inference.

Further, with evidence-based policymaking, correct inference of causal parameters also requires a correct stipulation of the government objective function into which the parameter estimates will be fed. That is, even with ideal first-stage policy randomization, one cannot avoid the need to make assumptions about the objective function determining governmental policy. To illustrate, suppose long-term policy is again endogenous, but with the government now applying a lower cutoff for deregulation, with $b^* = 0.50$. As shown in the dotted-dashed line in Figure 2, for each value of the causal parameter, the positive investment response at the start of the experiment will be smaller than in a the baseline economy featuring a higher cutoff (solid line). Intuitively, with a tougher deregulation standard, firms will assess a lower probability of deregulation post-experiment, resulting in a weaker response during the deregulation experiment. The failure to correctly account for a such change in government objective function parameters, here to a lower value of b^*, would cause causal parameter estimates to be biased downward.

The inference problems outlined above are readily understood analytically. The measured change in investment at the start of the experiment (the vertical axis in Figure 2), call it the deregulatory experiment 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) during and before the experiment are as shown in equations (18) and (20). Since these shadow values depend upon beliefs, as reflected by the presence of the $\beta(b)$ terms, it is apparent that any change in economic environment (e.g. F and b^*) shifting firm-level belief functions shifts the deregulatory experiment response function. Failure to account for such shifts results in incorrect causal parameter inference.

Consider next the special case of quadratic investment costs ($\nu = 2$). Here the deregulatory experiment response function is:

$$\begin{aligned} R(b) &= \left(\frac{1}{2\gamma}\right)\left[\frac{(r+\delta)(\overline{\pi}-\underline{\pi}) + \lambda_I[\beta(b)\overline{\pi} + (1-\beta(b))\underline{\pi} - \underline{\pi}]}{(r+\delta+\lambda_I)(r+\delta+\lambda_E)}\right] b \\ \Rightarrow \quad R'(b) &= \left(\frac{1}{2\gamma}\right)\left[\begin{array}{c} \frac{(r+\delta)(\overline{\pi}-\underline{\pi})+\lambda_I[\beta(b)\overline{\pi}+(1-\beta(b))\underline{\pi}-\underline{\pi}]}{(r+\delta+\lambda_I)(r+\delta+\lambda_E)} \\ +\beta'(b)\frac{\lambda_I(\overline{\pi}-\underline{\pi})b}{(r+\delta+\lambda_I)(r+\delta+\lambda_E)} \end{array}\right]. \end{aligned} \tag{24}$$

The preceding equations illustrate that a correct stipulation of the belief function $\beta(\cdot)$ is required to correctly infer the true causal parameter b based upon a measured experimental response. In turn, a correct stipulation of the belief function requires a correct stipulation of priors over the causal parameters (F) to be estimated, as well as the parameter of the government objective function (b^*).

Inspection of the derivative of the deregulatory investment response function reveals that endogenously heterogeneous beliefs serves to attenuate (amplify) investment response heterogeneity across the different possible b values under a Pro-Environment (Pro-Investment) government. Intuitively, a firm experiencing a high draw of b anticipates that a Pro-Environment (Pro-Investment) government will be less (more) likely to extend the experimental deregulation into the Implementation Stage, resulting in a weaker (stronger) experimental response.

Finally, comparing deregulation response functions across economies with endogenous versus exogenous policies, we find:

$$\underset{Endogenous}{R(b)} - \underset{Exogenous}{R(b)} = \left(\frac{\lambda_I}{2\gamma}\right)\left[\frac{\beta(b)\overline{\pi} + (1-\beta(b))\underline{\pi} - \pi_I^{EX}}{(r+\delta+\lambda_I)(r+\delta+\lambda_E)}\right] b. \tag{25}$$

It follows from equation (25) that there is necessarily a wedge between the experiment response functions across the Exogenous Policy and Endogenous Policy economies. In the former economy, post-experiment policy is determinate. In the latter economy, post-experiment policy is random, reflecting the uncertain outcome of econometric inference.

The following proposition summarizes the results of this subsection.

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.*

3.2 Observer Effects in NPEs

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

To this end, consider two economies facing the same type of deregulation NPE as discussed in the preceding subsection. In addition, assume both governments have the ability to choose long-term policy, $\pi_I \in \{\underline{\pi}, \overline{\pi}\}$, at their discretion. The only difference between the two economies is their respective governmental perceptions regarding the credibility of experimental evidence. In Economy NC, the government views experimental evidence as non-credible. In Economy C, the government views experimental evidence as credible.

Assume now that observation and measurement of firm behavior is either technologically feasible or not, 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 NC government ignores the evidence and implements the policy that is optimal given prior beliefs, call it π_{IP}^*. And if no observation occurs, the government has no choice but to rely on its priors, so it again implements π_{IP}^*. That is, firms in Economy NC will base their investments on the same conjectured long-term policy regardless of whether or not econometric 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. If they are not being observed, then the firms in Economy C anticipate that the government will need to rely on priors, so that $\pi_I = \pi^*_{IP}$. In contrast, under observation the firms anticipate implementation of the random policy $\pi^*_I(\mathbf{b})$. It follows that in Economy C, the jump in the shadow value of capital, and hence investment, at the onset of the experimental 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:

$$\underbrace{\left[\overline{q}^E(x,b) - \underline{q}^P(x,b)\right]}_{\text{Observed}} = \underbrace{\left[\overline{q}^E(x,b) - \underline{q}^P(x,b)\right]}_{\text{Not Observed}} + \underbrace{\frac{\lambda_I[\beta(b)\overline{\pi} + (1-\beta(b))\underline{\pi} - \pi^*_{IP}]b}{(r+\delta+\lambda_I)(r+\delta+\lambda_E)}}_{\text{Hawthorne Effect}}. \tag{26}$$

The preceding equation captures the fact that the act of econometric observation changes the distribution of the policy variable long-term, which will change shadow values and investment.

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.*

Figure 3 returns to the same parameter values as used in the preceding subsection's example (Figure 2). The figure now plots the experimental response function in our 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 (here $\pi^*_{IP} = \overline{\pi}$). 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 investment response to the experimental deregulation is attenuated. After all, firms are only certain of deregulation long-term if they are not observed.

4 Randomized Controlled Trials

This section considers econometric inference in the context of randomized controlled trials (RCTs) in which firms are randomly assigned to treatment and control groups during an experimental period. By construction, the RCTs satisfy Assumption 1 and are 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)).

4.1 Inference in RCTs

To fix ideas, we consider the following RCT. 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 from reverting back to its initial value, so $\pi_I = \overline{\pi}$. In the Endogenous Policy Economy, the government sets policy optimally given the evidence supplied by its econometrician.

For the purpose of the numerical example, 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 remaining parameter values are as follows: $r = .05$; $\delta = .10$; $\mu = 0$; $\lambda_E = 1$; $\lambda_I = 1$; $\overline{\pi} = 1$; $\underline{\pi} = 0$; $x = 1$; $\gamma = 1$; and $\nu = 4$.

Figure 4 illustrates RCT results. On the horizontal axis is the true value of the unknown causal 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 ($b^* = .85$) 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 where inference informs policy, the type of economy envisioned by some randomization advocates. The solid line in Figure 4 reflects 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. Apparently, the naïve econometrician will understate the true value of the causal parameter b. For example, suppose $b = 1$, resulting in an observed control-treatment investment difference of 100. Incorrectly, 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 non-representative if predicated upon a discretionary RCT, perhaps due to a government having private knowledge that the imposition of a novel regulation will be relatively harmless given the technological structure of its economy. Assumption 1 rules out this type of selection bias. 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 that with 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 causal parameter inference in the RCT is contingent upon a correct stipulation of the deep parameters of the government objective function into which the econometric estimates will be fed. To see this, suppose that the government were to adopt a higher threshold for deregulation. Then the control-treatment investment difference changes from the solid line to the dotted line. Thus, biased inference would result if an incorrect conjecture were to be made about the governmental objective function.

4.2 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 now, let us assume that observation may not be feasible, allowing us to assess whether the act of observation changes behavior.

The results of this exercise are shown 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. So why then does observation change behavior?

What differs between the observation and non-observation states 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 endogenously 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 here 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. More importantly, the act of observation changes the key test statistic here, the control-treatment investment difference. The next subsection sets out to understand why.

4.3 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:

$$\overline{q}^E(x,b) - \underline{q}^E(x,b) = \frac{(\overline{\pi} - \underline{\pi})b}{(r + \delta + \lambda_I)} \tag{27}$$

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. Intuitively, 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 policy expectations. That is, post-experiment policy expectations are necessarily equalized across treatment and control groups. Since post-experiment 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. Indeed, the difference between $\overline{q}^E$ and $\underline{q}^E$ shown in equation (27) is just the present value of a claim to the flow of excess pollution benefits $(\overline{\pi} - \underline{\pi})b$ accruing to the deregulated group during the Experiment Stage.

But recall, in the preceding subsection (Figure 5), 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 the econometrician, not the latent control-treatment shadow value difference (equation (27)). It is the behavior of this variable that we must understand.

To this end, let Δ denote the difference between control and treatment group investment. From equation (18) we have:

$$\begin{aligned}
\Delta(x,b) &= \underbrace{i^*[\overline{q}^E(x,b)]}_{Control} - \underbrace{i^*[\underline{q}^E(x,b)]}_{Treatment} \qquad (28) \\
&= \left(\frac{\nu-1}{\nu\gamma}\right)^{\nu-1}\left[\left(\underline{q}^E(x,b)+\frac{(\overline{\pi}-\underline{\pi})b}{(r+\delta+\lambda_I)}\right)^{\nu-1} - \left(\underline{q}^E(x,b)\right)^{\nu-1}\right] \\
&= \left(\frac{\nu-1}{\nu\gamma}\right)^{\nu-1}\left[\begin{array}{c} \left[\frac{x}{r+\delta-\mu}+\left(\underline{\pi}+\lambda_I\left(\frac{\beta(b)\overline{\pi}+(1-\beta(b))\underline{\pi}}{r+\delta}\right)\right)\left(\frac{b}{r+\delta+\lambda_I}\right)+\frac{(\overline{\pi}-\underline{\pi})b}{(r+\delta+\lambda_I)}\right]^{\nu-1} \\ -\left[\frac{x}{r+\delta-\mu}+\left(\underline{\pi}+\lambda_I\left(\frac{\beta_H(b)\overline{\pi}+(1-\beta(b))\underline{\pi}}{r+\delta}\right)\right)\left(\frac{b}{r+\delta+\lambda_I}\right)\right]^{\nu-1}\end{array}\right].
\end{aligned}$$

A key point to note in equation (28) is that beliefs (β) regarding policy post-experiment influence the investment of both the treatment and control groups during the experiment. Since the act of observation influences beliefs regarding long-term policy, 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 shape of the belief function is itself determined by F and b^*.

Despite the presence of observer effects for both treatment and control groups, it might be hoped that these 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 (28), in terms of the measured outcome variable i, in contrast to the unmeasured shadow value of capital q, observation effects do not generally cancel. In fact, 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 treatment and control groups cancel. In particular, it follows from equation (28) that:

$$\nu = 2 \Longrightarrow \Delta(x,b) = \left(\frac{1}{2\gamma}\right)\frac{(\overline{\pi}-\underline{\pi})}{(r+\delta+\lambda_I)} \times b. \qquad (29)$$

Notice, if $\nu = 2$, the control-treatment investment difference is linear in the causal parameter b. More importantly, the test statistic is now invariant to expectations regarding the distribution of the policy variable post-experiment. Thus, in the special case of quadratic investment costs, the control-treatment investment difference is utterly uncontaminated by any form of ex post endogeneity bias.

How general is this result? In order to provide a more complete characterization of the circumstances under which differences and (difference in differences) derived from RCTs are immune from post-experiment expectations contamination, we consider now a broader class of cost functions, reflective of those considered in the literature. 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 in dynamic settings.

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 (equation (2)) the property of being invariant to k, it follows that the shadow value formulae derived above (Subsection 2.4) remain valid. Second, Abel and Eberly (1994) show that under such a cost function, investment is weakly monotone increasing in q. Further, 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. With 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 difference in differences) was just shown to be invariant to post-experiment policy variable expectations 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 a 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.

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)*

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

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 the same parametric assumptions as 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, which is set to 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 firms that face an endogenous government policy response to the experiment, as well as 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 $b \in [.60, .74]$. Inactivity introduces a non-monotonicity into 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 incorrectly conclude that $b = 0$.

The dotted line considers that in addition to there being a wedge between the buy and sell price of capital (dashed line), 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 $b = 0$.

5 Non-Monotonicities and the Impossibility of Inference

Up to this point we have analyzed Rational Expectations Equilibria in which the government is able to correctly infer causal parameters based on some set of econometric evidence at its disposal. However, as we show now, NPEs and RCTs may by themselves be insufficient for this purpose.

Consider first the problem of causal parameter inference in the context of NPEs. To this end, Figure 7A returns to our baseline NPE (Figure 2) which featured an experimental deregulation. However, we consider now that the Pro-Environment government is now an even less-inclined to deregulate long-term, adopting a lower cutoff value of $b^* = .45$. Suppose now that we conjecture an equilibrium in which the government is able to determine the true causal parameter value based on the econometric evidence. However, suppose the only econometric evidence available to the government is the experimental outcome variable, the investment increase at the start of the deregulation experiment. It is apparent from Figure 7A that the behavior of the empirical outcome variable is inconsistent with the conjecture of correct inference in all states of nature. After all, the outcome variable is non-monotone in b. Consequently, the observed outcome cannot be inverted to solve for the true causal effect parameter b.

The root cause of the non-monotonicity here, and the impossibility of inference, is endogenous belief heterogeneity. Formally, this argument follows directly from equation (24) which reveals that the experimental investment response will be linear increasing, presenting no barrier to inference, if the belief function β is constant. To see the argument less formally, consider that holding beliefs fixed, firms drawing a higher b value would be inclined to invest more during the deregulation experiment, given that they receive a pollution benefit flow of πb per unit of installed capital. However, these same firms will rationally assign a lower probability to the Pro-Environment government continuing with deregulation long-term. This belief effect would curtail their investment incentive. Apparently, this expectations effect can be strong enough such that there is a region of b values, here from 0.80 to 0.90, where the investment response is actually decreasing in b. Thus, if the observed investment reaction were to fall into the region 0.80 to 1, the econometrician would not be able to infer the true value of the causal parameter.

Consider next the feasibility of identification of causal parameters based upon RCTs. To this end, Figure 7B returns to the same parameters assumed in our baseline RCT (Figure 4) but now assumes the macroeconomic profit factor x is close to zero. Suppose now that we conjecture an equilibrium in which the government is able to determine the true causal parameter value based on the RCT. However, suppose the only econometric evidence available to the government is the experimental outcome variable, the control-treatment group investment difference. As shown in Figure 7B, the outcome variable is here non-monotone in the underlying parameter. Absent other information, the government would not be able to invert the outcome variable to solve for the causal parameter.

6 Evidence-Based Policy-Making: Examples

This section illustrates that across a broad range of policy settings, not just those concerning financial markets, there is an increasing trend towards relying upon empirical evidence derived from randomizations in setting government policy variables, and where the anticipated time-path of those same policy variables can be expected to alter what is empirically measured during the observation window.

In 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 SEC, and the data 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. Of course, expectations regarding future short-sale regulations would have been reflected in the pricing and trading of all securities during the policy experiment. 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."

Consider next the relevance of our analysis for empirical evidence regarding agency and corporate finance theory. 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 would have been 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 ratchet-effects depend upon expectations regarding the future performance target, itself likely shaped by the experiment outcomes. Further, expectations regarding the relative weight placed on future discretionary payments will influence the willingness of tax collectors to accumulate reputational capital through honest dealing during the experiment.

Consider next environmental economics. In conjunction with President Obama's landmark Clean Policy Plan, the EPA 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 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 development economics, randomized evidence-based policymaking 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, say randomized extension of micro-credit, 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 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."

In 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 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 the fiscal stimulus proposed by President Obama. 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 likely that Romer, in her policy-making role, was heavily influenced by her own econometric evidence, and discounted evidence coming from sources deemed to be less credible.

7 Conclusion

This paper illustrates an inherent tradeoff between the credibility of empirical estimates derived from randomization, and their practical utility in dynamic settings. 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.

These econometric challenges 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. First, it is apparent that the severity of bias 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 biases will be less severe. 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] Acharya, Viral and Paolo Volpin, 2009, Corporate governance externalities, *Review of Finance* 14 (1), 1-33.

[5] 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.

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

[7] 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.

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

[9] 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.

[10] Biais, Bruno and Paul Woolley, 2011, High frequency trading, Working Paper, Toulouse School of Economics.

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

[12] Bond, Philip, Itay Goldstein, and Edward Simpson Prescott, 2009, Market-based corrective actions, *Review of Financial Studies*, 782-820.

[13] Chassang, Sylvain, Gerard Padro i Miguel, and Erik Snowberg, 2012, Selective trials: A principal-agent approach to randomized controlled experiments, *American Economic Review* (102), 1279-1309.

[14] 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.

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

[16] 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.

[17] Gorton, Gary, 2010, *Slapped by the Invisible Hand*, Princeton University Press, Princeton, N.J.

[18] 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.

[19] Hart, Oliver, 2009, Regulation and Sarbanes-Oxley, *Journal of Accounting Research* 47 (2), 437-445.

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

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

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

[23] 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.

[24] 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.

[25] 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.

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

[27] 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.

[28] 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.

[29] Romer, Paul, 2016, The trouble with macroeconomics, forthcoming in *The American Economist.*

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

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

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

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

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

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

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

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

[38] 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:

$$rV^I(x,k) = \max_i (x + \pi_I b)k - \gamma i^{\nu/(\nu-1)} + \mu x V_x^I(x,k) + \frac{1}{2}\sigma^2 x^2 V_{xx}^I(x,k) + (i - \delta k)V_k^I(x,k). \quad (30)$$

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

$$V^I(x,k) = kq^I(x) + G^I(x). \quad (31)$$

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:

$$(r+\delta)q^I(x) = (x + \pi_I b) + \mu x q_x^I(x) + \frac{1}{2}\sigma^2 x^2 q_{xx}^I(x). \quad (32)$$

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

$$q^I(x) = x\zeta^I + \eta^I. \quad (33)$$

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

$$\begin{aligned} \zeta^I &= \zeta \equiv 1/(r + \delta - \mu). \\ \eta^I &= \frac{\pi_I b}{r+\delta} \end{aligned} \quad (34)$$

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

$$q^I(x) = \zeta x + \eta^I \Rightarrow i^*[q^I(x)]q^I(x) - \gamma[i^*(q^I(x))]^{\nu/(\nu-1)} = (\zeta x + \eta^I)^\nu \Gamma = \sum_{h=0}^{\nu} \phi_h^I x^h \quad (35)$$

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:

$$\widehat{r}g_h(x) = \mu x g_h'(x) + \frac{1}{2}\sigma^2 x^2 g_h''(x) + \phi_h x^h \tag{36}$$

We conjecture this value function takes the form:

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

Substituting the conjectured solution back into equation (36) 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:

$$G^I(x) = \sum_{h=0}^{\nu} \phi_h^I \omega_h^I x^h \tag{37}$$

$$\text{with } \phi_h^I \equiv \binom{\nu}{h} (\eta^I)^{\nu-h} \zeta^h \Gamma.$$

$$\omega_h^I \equiv \frac{1}{r - \mu h - \frac{1}{2}\sigma^2 h(h-1)}.$$

Experiment Stage

The HJB equation for the Experiment Stage is:

$$\begin{aligned} rV^E(x,k) &= \max_i \quad (x + \pi_E b)k - \gamma i^{\nu/(\nu-1)} + \mu x V_x^E(x,k) + \frac{1}{2}\sigma^2 x^2 V_{xx}^E(x,k) \\ &\quad + (i - \delta k)V_k^E(x,k) + \lambda_I \beta \left[\overline{V}^I(x,k) - V^E(x,k)\right] + \lambda_I(1-\beta)\left[\underline{V}^I(x,k) - V^E(x,k)\right]. \end{aligned} \tag{38}$$

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

$$V^E(x,k) = kq^E(x) + G^E(x). \tag{39}$$

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:

$$(r+\delta+\lambda_I)q^E(x) = (x+\pi_E b) + \mu x q_x^E(x) + \frac{1}{2}\sigma^2 x^2 q_{xx}^E(x) + \lambda_I \left[\frac{x}{r+\delta-\mu} + \frac{(\beta\overline{\pi} + (1-\beta)\underline{\pi})b}{r+\delta}\right]. \tag{40}$$

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

$$q^E(x) = x\zeta^E + \eta^E. \tag{41}$$

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

$$\zeta^E = \zeta \equiv 1/(r+\delta-\mu). \tag{42}$$

$$\eta^E = \left[\frac{(r+\delta)\,\pi_E b + \lambda_I(\beta\overline{\pi} + (1-\beta)\underline{\pi})b}{(r+\delta)\,(r+\delta+\lambda_I)}\right].$$

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:

$$(r+\lambda_I)\,G^E(x) = \mu x G_x^E(x) + \frac{1}{2}\sigma^2 x^2 G_{xx}^E(x) + \lambda_I\left[\beta\overline{G}^I(x) + (1-\beta)\underline{G}^I(x)\right] + \sum_{h=0}^{\nu}\left[\binom{\nu}{h}\left(\eta^E\right)^{\nu-h}\zeta^h\Gamma\right]x^h. \tag{43}$$

This ODE can be rewritten as:

$$(r+\lambda_I)\,G^E(x)=\mu x G^E_x(x)+\frac{1}{2}\sigma^2x^2G^E_{xx}(x)+\sum_{h=0}^{\nu}\phi^E_h x^h. \tag{44}$$

where

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

From the Growth Option Lemma it follows:

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

Pre-Experiment Stage

The HJB equation is:

$$\begin{aligned}
rV^P(x,k) \;=\; \max_i \quad & (x+\pi_P b)k-\gamma i^{\nu/(\nu-1)}+\mu x V^P_x(x,k)+\frac{1}{2}\sigma^2x^2V^P_{xx}(x,k) \\
&+(i-\delta k)V^P_k(x,k)+\lambda_E\theta\left[\overline{V}^E(x,k)-V^P(x,k)\right]+\lambda_E(1-\theta)\left[\underline{V}^E(x,k)-V^P(x,k)\right].
\end{aligned} \tag{45}$$

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

$$V^P(x,k)=kq^P(x)+G^P(x). \tag{46}$$

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:

$$(r+\delta+\lambda_E)q^P(x)=(x+\pi_P b)+\mu x q^P_x(x)+\frac{1}{2}\sigma^2x^2q^P_{xx}(x)+\lambda_E\theta\overline{q}^E(x)+\lambda_E(1-\theta)\underline{q}^E(x). \tag{47}$$

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:

$$(r+\lambda_E)G^P(x) = \mu x G^P_x(x) + \frac{1}{2}\sigma^2 x^2 G^P_{xx}(x) + \sum_{h=0}^{\nu}\left[\binom{\nu}{h}(\eta^P)^{\nu-h}\zeta^h\Gamma\right]x^h + \lambda_E\left[\theta\overline{G}^E(x) + (1-\theta)\underline{G}^E(x)\right]. \tag{48}$$

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

$$(r+\lambda_E)G^P(x) = \mu x G^P_x(x) + \frac{1}{2}\sigma^2 x^2 G^P_{xx}(x) + \sum_{h=0}^{\nu}\phi^P_h x^h. \tag{49}$$

with

$$\begin{aligned}
\phi^P_h &\equiv \binom{\nu}{h}(\eta^P)^{\nu-h}\zeta^h\Gamma + \lambda_E\omega^E_h\left[\theta\overline{\phi}^E_h + (1-\theta)\underline{\phi}^E_h\right] \\
\overline{\phi}^E_h &\equiv \left[\lambda_I\left(\beta\overline{\phi}^I_h + (1-\beta)\underline{\phi}^I_h\right)\omega^I_h\right] + \binom{\nu}{h}\left(\overline{\eta}^E\right)^{\nu-h}\zeta^h\Gamma \\
\underline{\phi}^E_h &\equiv \left[\lambda_I\left(\beta\overline{\phi}^I_h + (1-\beta)\underline{\phi}^I_h\right)\omega^I_h\right] + \binom{\nu}{h}\left(\underline{\eta}^E\right)^{\nu-h}\zeta^h\Gamma.
\end{aligned} \tag{50}$$

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:

$$\begin{aligned}
G^P &= \sum_{h=0}^{\nu}\phi^P_h\omega^P_h x^h \\
\omega^P_h &\equiv \frac{1}{r+\lambda_E-\mu h-\frac{1}{2}\sigma^2 h(h-1)}.
\end{aligned} \tag{51}$$

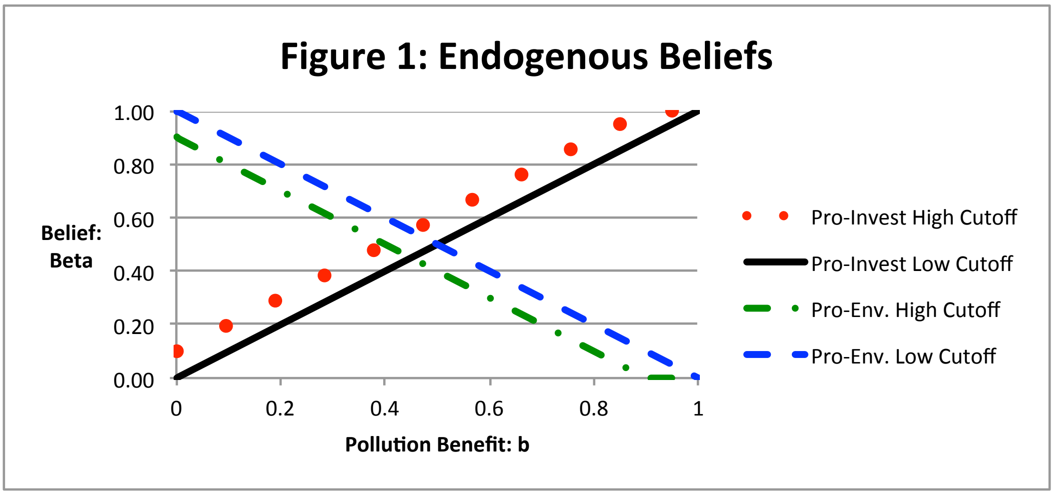
Figure 1: Endogenous Beliefs
1.00
0.80
0.60
0.40
0.20
0.00
Belief: Beta
0
0.2
0.4
0.6
0.8
1
Pollution Benefit: b
Pro-Invest High Cutoff
Pro-Invest Low Cutoff
Pro-Env. High Cutoff
Pro-Env. Low Cutoff

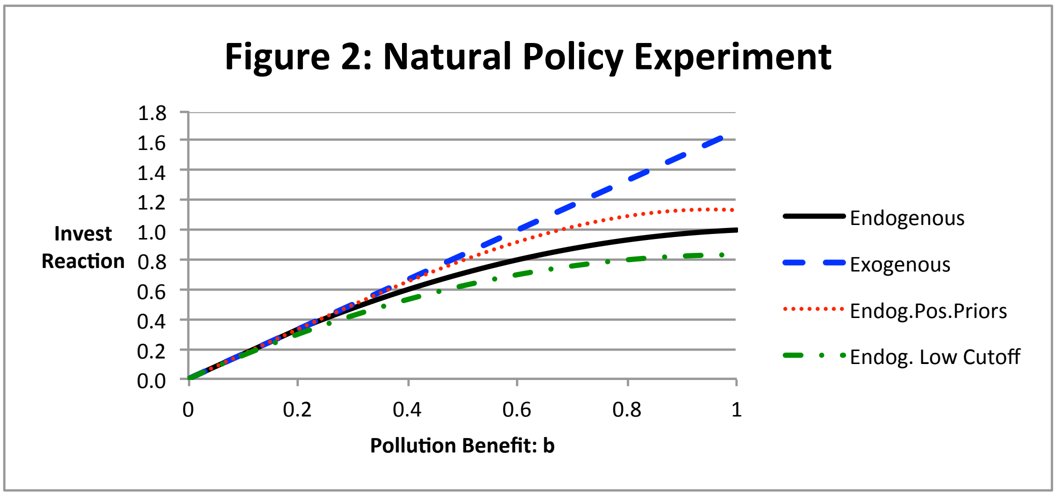
Figure 2: Natural Policy Experiment
1.8
1.6
1.4
1.2
1.0
0.8
0.6
0.4
0.2
0.0
Invest Reaction
0
0.2
0.4
0.6
0.8
1
Pollution Benefit: b
Endogenous
Exogenous
Endog.Pos.Priors
Endog. Low Cutoff

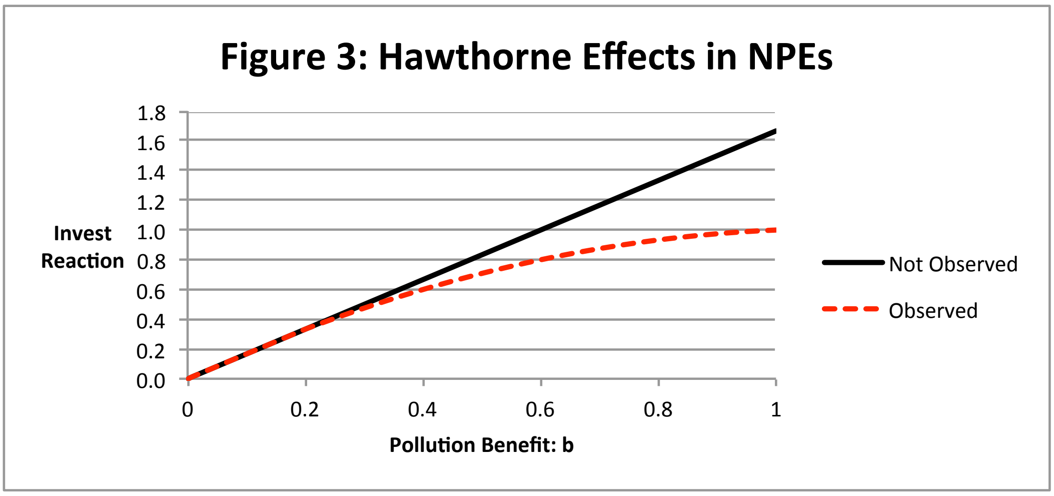
Figure 3: Hawthorne Effects in NPEs
1.8
1.6
1.4
1.2
1.0
0.8
0.6
0.4
0.2
0.0
Invest Reaction
0
0.2
0.4
0.6
0.8
1
Pollution Benefit: b
Not Observed
Observed

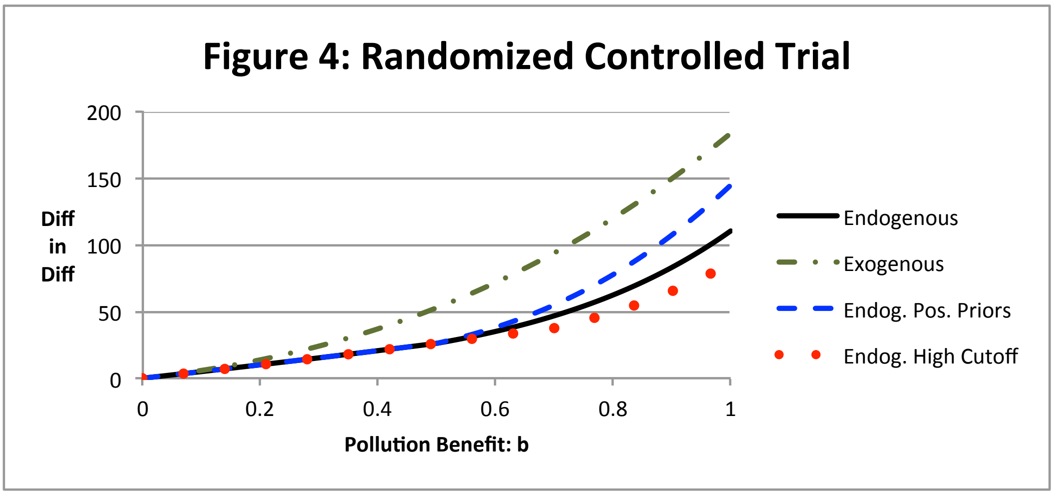
Figure 4: Randomized Controlled Trial
200
150
100
50
0
Diff
in
Diff
0
0.2
0.4
0.6
0.8
1
Pollution Benefit: b
Endogenous
Exogenous
Endog. Pos. Priors
Endog. High Cutoff

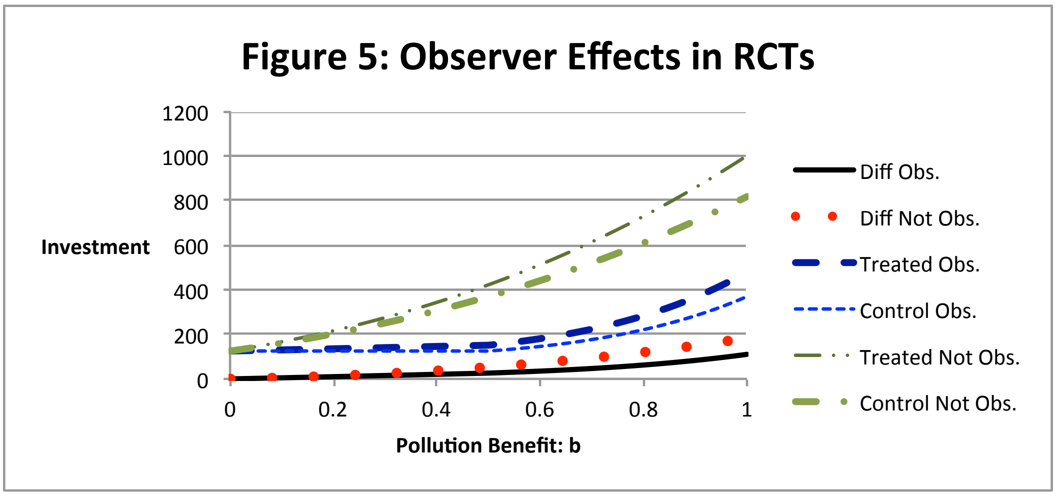
Figure 5: Observer Effects in RCTs
1200
1000
800
600
400
200
0
Investment
0
0.2
0.4
0.6
0.8
1
Pollution Benefit: b
Diff Obs.
Diff Not Obs.
Treated Obs.
Control Obs.
Treated Not Obs.
Control Not Obs.

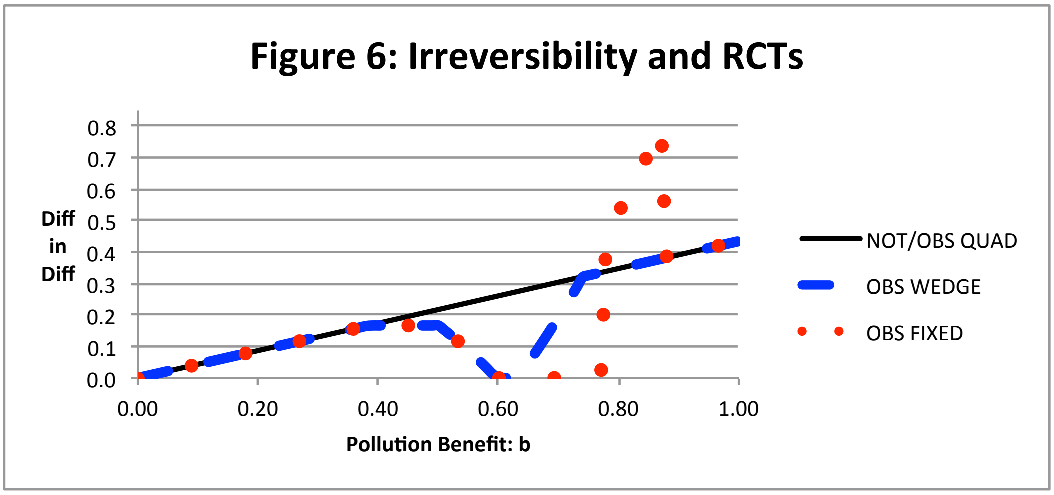
Figure 6: Irreversibility and RCTs
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0.0
Diff
in
Diff
0.00
0.20
0.40
0.60
0.80
1.00
Pollution Benefit: b
NOT/OBS QUAD
OBS WEDGE
OBS FIXED

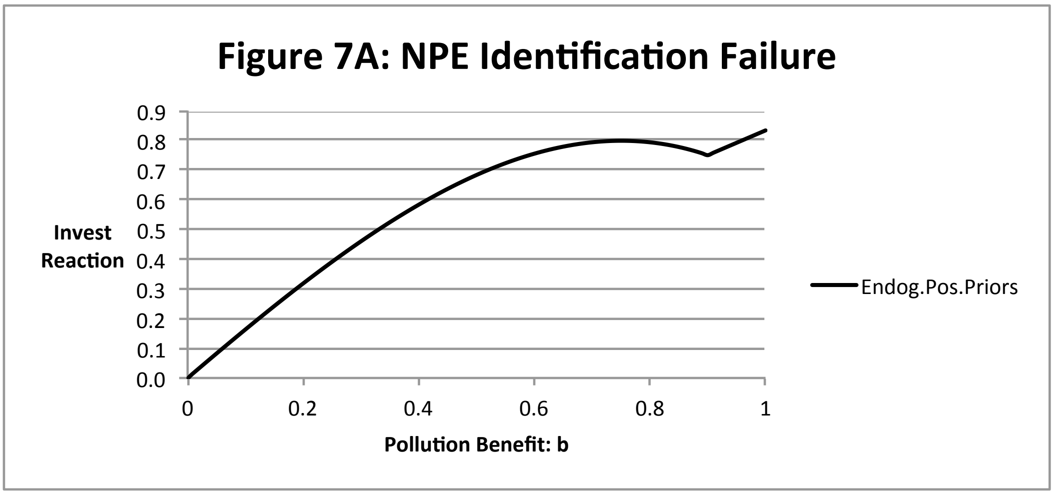
Figure 7A: NPE Identification Failure
Invest Reaction
0.0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
0
0.2
0.4
0.6
0.8
1
Pollution Benefit: b
Endog.Pos.Priors

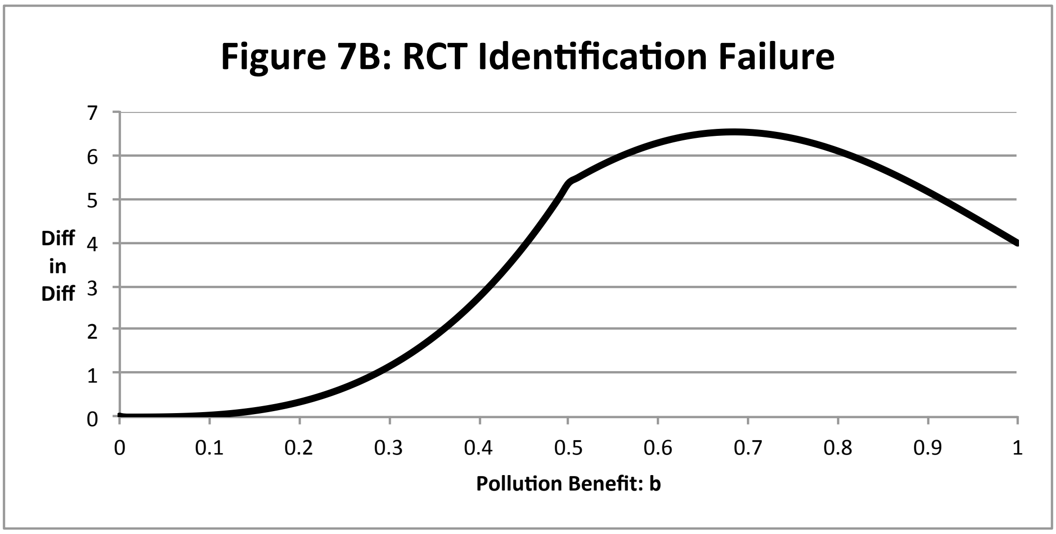
Figure 7B: RCT Identification Failure
Diff in Diff
0
1
2
3
4
5
6
7
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
1
Pollution Benefit: b