

# DISCUSSION PAPER SERIES

DP11361

## **THE PARADOX OF POLICY-RELEVANT NATURAL EXPERIMENTS**

Gilles Chemla and Christopher Hennessy

**DEVELOPMENT ECONOMICS,  
FINANCIAL ECONOMICS, INDUSTRIAL  
ORGANIZATION and PUBLIC  
ECONOMICS**



# THE PARADOX OF POLICY-RELEVANT NATURAL EXPERIMENTS

*Gilles Chemla and Christopher Hennessy*

Discussion Paper DP11361  
Published 28 June 2016  
Submitted 08 December 2017

Centre for Economic Policy Research  
33 Great Sutton Street, London EC1V 0DX, UK  
Tel: +44 (0)20 7183 8801  
[www.cepr.org](http://www.cepr.org)

This Discussion Paper is issued under the auspices of the Centre's research programme in **DEVELOPMENT ECONOMICS, FINANCIAL ECONOMICS, INDUSTRIAL ORGANIZATION and PUBLIC ECONOMICS**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Gilles Chemla and Christopher Hennessy

# THE PARADOX OF POLICY-RELEVANT NATURAL EXPERIMENTS

## Abstract

We examine robustness of evidence derived from ideal randomizations applied to atomistic subjects in dynamic settings. Paradoxically, once experimental evidence is viewed as sufficiently clean to use, it then becomes contaminated by ex post endogeneity: Measured responses depend upon priors and the objective function into which evidence is fed. Moreover, agents policy beliefs become endogenously correlated with their causal parameters, clouding inference. Finally, treatment-control differences are contaminated absent quadratic adjustment costs. Constructively, we illustrate how inference can be corrected accounting for feedback and highlight factors mitigating contamination.

JEL Classification: N/A

Keywords: randomized controlled trials, natural policy experiments, firms, investment, pollution, government

Gilles Chemla - g.chemla@imperial.ac.uk  
*Imperial College Business School, CNRS, and CEPR and CEPR*

Christopher Hennessy - chennessy@london.edu  
*London Business School, CEPR, and ECGI and CEPR*

## Acknowledgements

We thank seminar participants at CEPR ESSET (Experimental), Stanford, U.C. Berkeley, Columbia, CMU, MIT, LBS BC, INSEAD, IESE, FTG LSE, Imperial, Baruch, Copenhagen, Washington-Seattle, Miami, Paris-Dauphine, Zurich, Arison, and Nova-Lisbon. Hennessy acknowledges funding from the ERC.

# THE PARADOX OF POLICY-RELEVANT NATURAL EXPERIMENTS\*

Gilles Chemla

Imperial College, DRM/CNRS, and CEPR.

Christopher A. Hennessy

LBS, CEPR, and ECGI

July 2017

## Abstract

We examine robustness of evidence derived from ideal randomizations applied to atomistic subjects in dynamic settings. Paradoxically, once experimental evidence is viewed as sufficiently clean to use, it then becomes contaminated by *ex post endogeneity*: Measured responses depend upon priors and the objective function into which evidence is fed. Moreover, agents' policy beliefs become endogenously correlated with their causal parameters, clouding inference. Finally, treatment-control differences are contaminated absent quadratic adjustment costs. Constructively, we illustrate how inference can be corrected accounting for feedback and highlight factors mitigating contamination.

---

\*We thank seminar participants at CEPR ESSET (Experimental), Stanford, U.C. Berkeley, Columbia, CMU, MIT, LBS BC, INSEAD, IESE, FTG LSE, Imperial, Baruch, Copenhagen, Washington-Seattle, Miami, Paris-Dauphine, Zurich, Arison, and Nova-Lisbon. Hennessy acknowledges funding from the ERC.

# 1 Introduction

In their influential textbook, *Mostly Harmless Econometrics*, Angrist and Pischke (2009) argue empirical evidence derived from exogenous shocks and random assignment 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.”

While undoubtedly constituting a methodological advance, the “credibility revolution” has also faced some criticism for at times seeming to elevate identification over economic relevance (see e.g. Rust (2010) and Keane (2010)). However, in response to such criticisms, applied micro-econometricians have redoubled their efforts to demonstrate policy relevance. In fact, the perceived credibility of evidence derived from randomization has recently led to calls for heavy/exclusive reliance on such evidence in setting policy. For example, Dhaliwal and Tulloch (2015) note the existence of an “increasing trend towards considering rigorous evidence while making policy decisions.” Consistent with this view, the mission of J-PAL is to promote “evidence-based policymaking.” Similarly, Greenstone (2009) calls for government agencies “to move toward a culture of persistent regulatory experimentation” in which randomized regulations are sunsetted so that impact analysis can inform the next regulatory decision. Duflo (2004) argues, “Creating a culture in which rigorous randomized evaluations are promoted, encouraged, and financed has the potential to revolutionize social policy during the 21st century, just as randomized trials revolutionized medicine during the 20th.”

In this paper, we view evidence derived from seemingly-ideal randomizations through the lens of rational expectations (Muth (1961)). In so doing, our goal is not to add to the stock of debating points accumulated on either side of the structural versus reduced-form debate (surveyed below), but rather to develop an analytical framework that refines our understanding what randomization ultimately can and cannot deliver in the policy arena, an arena in which the methodological tool-kit has enjoyed increasing popularity on the basis of the perceived simplicity and credibility.

Our model reveals an internal inconsistency in the natural experiment research program specifically as it relates to causal parameter inference in dynamic settings. A simple example illustrates. We begin by noting that the two primary objectives in contemporary applied micro-econometric work are to convince the audience of clean identification and policy relevance. Suppose then that after exhaustive debate the econometrician is ultimately able to convince her audience that her identification strategy is clean in the sense that she can offer a compelling demonstration based upon her granular knowledge of institutional details that nature itself forced an exogenous change in government policy and/or randomly assigned agents to treatment and control groups. The econometrician is next challenged on policy-relevance. Suppose that here too, she can rise to the challenge,

demonstrating direct policy-relevance, e.g. the experimental outcomes are actually being utilized by the government in considering the scale-up of the initial pilot experiment.

What has gone unnoticed here is that there is a contradiction between the econometrician’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 econometrician 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, they will have rationally changed their behavior during the experiment anticipating the influence of the econometric evidence. That is, the expectation of endogenous evidence-based policymaking after the experiment will change what the econometrician measures during the experiment. Importantly, as is the case in our model, this is true even if one considers an ideal “large-sample” setting where each agent is measure zero and has no strategic motive to distort behavior with the goal of influencing subsequent policy.

What implications does this have for causal inference? To address this question formally, we consider the following parable economy. At each point in time, atomistic firms operating across a finite number of industries make investment decisions in light of current and expected future regulation. Tight regulation reduces the flow of unobservable private benefits accruing to owner-managers. The industry-specific private benefits are unobservable, being i.i.d. draws from a known probability distribution. Econometricians and the government want to infer the magnitude of private benefits in different sectors, since private benefits determine the key causal effect parameters in this economy: the effect of tighter regulation on long-term industry-level investment.

Fortunately for the econometricians in the parable 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 below), all firms are subjected to a common exogenous shock to regulations during a randomly-timed Experiment Stage. In a *Randomized Controlled Trial* (RCT below), a fraction of firms randomly face regulation during the Experiment Stage and the remaining firms do not. In the NPE setting, econometricians attempt to infer sector-specific private benefits based upon the change in investment induced by the experimental regulation shock. In the RCT setting, inference instead relies upon the observed difference between treatment and control group investment.

As we show, feedback from experimental evidence to the probability distribution of the policy variable post-experiment contaminates the formerly-clean evidence. In particular, evidence that is used to inform policy decisions (“policy-relevant evidence”) is contaminated by what we term *ex post endogeneity*. And this is true even though, by construction, the individual firms in our economy are measure zero and have no ability to influence empirical test statistics or policy decisions. However, the problem of ex post endogeneity vanishes if the government is powerless to change future policy.

Similarly, the contamination vanishes if the government does not view the experimental evidence as credible and ignores it. We thus have the following paradoxical situation here: Seemingly-clean experimental evidence is uncontaminated only if the government is unable or unwilling to use it.

We move well beyond illustrating this paradox, describing five novel challenges to causal parameter inference arising from ex post endogeneity. First, rather than being stand-alone objects that are “not heavily model-dependent,” policy-relevant experimental evidence must be interpreted in light of the governmental objective function into which the evidence will be fed. That is, the fact that one has initially observed an ideal policy randomization drawn independently of the government objective function does not eliminate the need to make assumptions about the government’s objective function. After all, the government’s objective function will still influence the distribution of the policy variable after the experiment, thus influencing measured responses of forward-looking agents during the experiment. Effectively, even with ideal randomization, the contaminating effects of policy endogeneity have 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 governing these same causal parameters. Intuitively, since agents rationally anticipate the government setting its post-experiment policy based upon the inferred values of causal parameters, prior beliefs regarding the probability distribution of these causal parameters influence agent beliefs and behavior. Consequently, an incorrect stipulation of prior beliefs regarding the probability distribution governing causal parameters will lead to incorrect inference. Phrased differently, external validity (extrapolation) of an experimental result requires the strong assumption that, across alternative settings, agents hold the same prior beliefs regarding the probability distribution governing causal parameters.

Third, the ex post endogeneity problem described above generates observer effects: the act of observation by econometricians changes the measured responses of agents in both the treatment group (Hawthorne Effect) and control group (John Henry Effect).<sup>1</sup> The behavioral and organization literatures have postulated a range of rationales for observer effects such as self-consciousness, approval-seeking, spite, or a desire to influence study outcomes. However, by construction, our model abstracts from each of these stories since firms are rational, measure zero, and anonymous.

Fourth, the observer effects described in the preceding paragraph are unequal across treatment and control groups in RCTs unless the underlying stock variable accumulation technology satisfies strong functional assumptions: zero fixed costs, equality of buy and sell prices, and quadratic adjustment costs. As shown, if these functional form assumptions are not satisfied, treatment-

---

<sup>1</sup>See Levitt and List (2011) and Zwane et al. (2011).

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, if agents are endowed with heterogeneous causal effect parameters, they will have 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. Failure to take this into account will lead to the confounding of beliefs and causal parameters, resulting in faulty inference. Worse still, experimental moments may 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 key feature of the model is that choices made during the experiment have implications for agents' future payoffs, with these payoffs being influenced by future government policy. Thus, the arguments here are most directly relevant to settings in which government policies influence the accumulation of stock variables at the firm-level, e.g. capital, patents, cash hoarding, debt, and employees. Similarly, our arguments apply to a broad range of dynamic household decision margins such as savings, portfolio allocation, lifetime labor supply, offspring, and human capital accumulation.

The type of evidence-policy feedback at the heart of our model is common. Evidence regarding responses to environmental regulations informed EPA decisions under the Obama Administration. The U.S. tax changes of the mid-1980's were informed by an exhaustive empirical literature analyzing responses to the tax code overhaul that took place in the early 1980's. Empirical analyses of corporate responses to Sarbanes-Oxley have informed accounting and governance regulations. More recently, as described by Spatt (2011), the SEC randomly relaxed the up-tick rule on a sample of stocks and then, based upon the observed responses, repealed the rule for all stocks.

We do not claim the biases we illustrate will be large in every dynamic setting. Rather, we simply argue biases will become larger the tighter the nexus between estimation and policy-setting. But at this juncture we hasten to point out that a tight estimation-policy nexus is actively promoted by many. Moreover, credibly informing policy is a key goal of much empirical work.

The issues raised here 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 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 show that *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 argument 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 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 paradox we illustrate, upon which Lucas (1976) is silent. These differences notwithstanding, the present paper borrows from Lucas (and Muth (1961)) the idea of viewing empirical evidence, here policy-relevant experiments, through the prism of rational expectations.

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 analysis 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 their assigned 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 deliberately shut off in our model. Acemoglu also argues 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. Despite these differences, both papers share the goal of filtering experimental evidence through the lens of a model.

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, rendering causal parameter inference impossible. We conclude with some practical guidance regarding the types of experimental evidence that one should look for in light of our analysis.

## 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 policy variable. The model itself 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:

$$dk_t = (i_t - \delta k_t)dt. \quad (1)$$

The variable  $i$  denotes gross investment and  $\delta \geq 0$  is the depreciation rate. Firms invest optimally each instant.

The investment cost function is common to all firms and is common knowledge to all agents, including the government. It takes the following simple form:

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

To ensure the optimal instantaneous control policy is unique, we assume the cost function parameters satisfy  $\gamma > 0$  and  $\nu > 1$ . Here we will obtain a simple closed-form expression 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  $\nu$ , as in Abel and Eberly (1997). The literature on dynamic accumulation problems, e.g. Abel and Eberly (1994, 1997), has also considered fixed costs and irreversibilities. We have abstracted from such complexities here, but will discuss their implications as model extensions.

Total profits, inclusive of non-monetary private benefits and costs, cannot be observed by outsiders, including the government. Date  $t$  profits are:

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

In the model, regulation impairs the ability of manager-owners to capture non-monetary private benefits of ownership and control, creating disincentives for business growth, with the government seeking to infer the magnitude of the disincentive effect. For example, the utility derived from

running a business is reduced by the time and effort costs associated with complying with regulations. As a second example, complying with regulations may force firms to disclose information they would prefer to keep private. As a final example, regulations often force businesses to adopt practices they disagree with.

In the profit equation (3), the variable  $b$  measures the flow of private benefits the firm’s manager-owner would receive in the absence of regulation. The term  $\pi_t$  represents the percentage of potential private benefits captured by the manager-owner accounting for regulation. The interval of possible regulation levels is  $\pi_t \in [\underline{\pi}, \bar{\pi}]$ , where  $0 \leq \underline{\pi} < \bar{\pi} \leq 1$ . For example, one can think of  $\underline{\pi} = 0$  and  $\bar{\pi} = 1$ , in which case  $b$  measures private benefits that are lost when the government departs from laissez-faire and imposes a regulation that cuts off the flow of private benefits.

Private benefits are assumed to be industry-specific, and there are  $M \geq 2$  industries in the economy. The key ingredient here is that within each industry, firm technologies are (perfectly) correlated. Therefore, a firm’s own technology is informative about the firms in its industry. This will help individual firms to better forecast endogenous government policy decisions post-experiment.

The industry-specific private benefit parameters are i.i.d. draws at date 0. The private 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  $\tilde{\mathbf{b}}$ , is denoted  $\mathbf{b}$ . Econometricians and the government want to infer the realized vector of private benefit parameters  $\mathbf{b}$  since this will allow them to infer the magnitude of investment responses to changes in regulation. Below we speak of the cumulative distribution function  $F$  as capturing *prior beliefs* (over the unknowns here, the realized causal parameter vector  $\mathbf{b}$ ).

The stochastic profit factor  $x$  entering the profit equation (3) is a positive geometric Brownian motion with the following law of motion:

$$dx_t = \mu x_t dt + \sigma x_t dw_t. \quad (4)$$

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

All firms in the economy face the same stochastic profit factor  $x$ . This assumption is not necessary, but serves the purpose of approximating the type of macroeconomic shocks to which a real-world government might be expected to respond, opening up the potential for endogeneity bias. For example, one might be concerned about downward bias in causal parameter estimates if governments are more willing to impose regulations in good times (high  $x$ ). Anticipating, the exogenous policy randomizations we consider below are such that standard forms of (ex ante) endogeneity and selection bias will not be an issue—the experiments considered will be optically ideal.

## 2.2 Timing

It is convenient to split the model into three stages,  $S \in \{P, E, I\}$ . A firm faces constant regulation within a given stage, but varying regulation across stages. The regulation variable during stage  $S$  is denoted  $\pi_S$ . The Pre-Experiment Stage  $P$  is followed by the Experiment Stage  $E$  which is followed by the Implementation Stage  $I$ . During each stage, firms face a simple time-homogeneous instantaneous investment problem, so the setup and solution are essentially equivalent to a simple three period model.

During Stage  $P$ , all firms face the same economy-wide regulatory standard  $\pi_P \in [\underline{\pi}, \bar{\pi}]$ . This can be thought of as the initial endowed technology in the economy. An exogenous natural experiment will arrive at date  $\tilde{\tau}_E$ . This date is an independent random variable. The transition rate into the Experiment Stage is  $\lambda_E > 0$ . Thus, at any time prior to the transition, the expected remaining duration of Stage  $P$  is  $\lambda_E^{-1}$ .

During the Experiment Stage, Nature randomly assigns a fraction  $\theta$  of firms to “fully deregulated status” where  $\pi_E = \bar{\pi}$ , with the remaining firms assigned to “fully regulated status” where  $\pi_E = \underline{\pi}$ . Two types of experiments are considered. In a *Natural Policy Experiment* (NPE below), all firms face a common exogenous regulatory policy  $\pi_E \neq \pi_P$  during the Experiment Stage. In a *Randomized Controlled Trial* (RCT below), firms are randomly assigned to treatment and control groups, one fully regulated the other fully unregulated. In modeling an RCT, the parameter  $\theta \in (0, 1)$ . For example, in an RCT featuring an equal measure of firms in treatment and control groups,  $\theta = 1/2$ .

Unexpected experiments are captured by setting  $\lambda_E$  to an arbitrarily small number. This will have no effect on treatment and control group differences in RCTs. For NPEs, unexpected experiments simply feature relatively large investment changes at the start of the experiment. As shown below, such a level shift in the policy reaction function has no bearing on our analysis of ex post endogeneity and concomitant bias in inference.

The Implementation Stage  $I$  will arrive at date  $\tilde{\tau}_I$ . This date is an independent random variable given  $\tilde{\tau}_E$ . 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  $\lambda_I^{-1}$ .

At the very start of Stage  $I$ , econometricians have the opportunity to observe some experimental evidence.<sup>2</sup> As shown below, the causal parameter vector  $\mathbf{b}$  can potentially be correctly inferred from this experimental evidence, but only if the econometrician understands the subtle interplay between evidence, policy, and firm-level expectations.<sup>3</sup>

In an RCT, econometricians can look back and measure the difference between the investment of

---

<sup>2</sup>Letting firms observe the same evidence as econometricians has no effect.

<sup>3</sup>Observation during Stage  $E$  adds an uninteresting waiting phase where  $\pi_I$  is determined, with  $\pi_E$  still in effect.

regulated and unregulated firms, industry-by-industry. 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 stochastic 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 following assumption is satisfied by the stochastic regulatory policy process  $\tilde{\pi}_{jt}$  facing arbitrary firm  $j$  at any point in time  $t$  at which the experimental measurement may take place.

**Assumption 1 (Independence):**  $\tilde{\pi}_{jt} \perp \{\tilde{\mathbf{b}}, \tilde{x}_t\} \quad \forall \quad j \in \mathbb{J} \text{ and } t \in [0, \tilde{\tau}_I]$ .

By construction, independence of the regulatory policy process rules out the standard forms of endogeneity bias. First, random assignment rules out selection by firms or the government based on unobservables ( $\mathbf{b}$ ), e.g. firms 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 a novel regulation, say  $\pi_E = \underline{\pi}$ , because it knows that the true causal parameter configuration in its economy is benign. However, in the economy considered, the government does not know the parameter vector  $\mathbf{b}$ , and the policy variable is independent of the random vector  $\tilde{\mathbf{b}}$  at any point in time at which the experimental measurement may take place, specifically all times  $t \in [0, \tilde{\tau}_I]$ . 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 macroeconomic state ( $x_t$ ). In this way, the experiment has seemingly-ideal optics.

During the Implementation Stage, a long-term regulatory policy ( $\pi_I$ ) will be implemented permanently. In the Exogenous Policy Economy,  $\pi_I$  will be set to a technologically pre-determined value  $\pi_I^{EX}$  and the government is powerless to change it. In the Endogenous Policy Economy, the government will, with the help of its sophisticated econometrician, infer the true 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, the government in the Endogenous Policy Economy will implement a policy  $\pi_I^*$  to maximize an objective function  $\Theta$ . Specifically:

$$\pi_I^*(\mathbf{b}) \in \arg \max_{\pi \in [\underline{\pi}, \bar{\pi}]} \Theta(\pi, \mathbf{b}) \quad (5)$$

where

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

Notice, with such an objective function, the government will utilize a bang-bang decision rule: fully deregulate (regulate) if the average increase in private benefits resulting from deregulation is greater (less) than the critical threshold  $b^*$ . In other words, the government deregulates if the average cost regulation imposes on firms is deemed to be too large. Moreover, since, as shown below, long-term investment is increasing in private benefits, one can think of the government here as being willing to deregulate if doing so will bring about a sufficiently large increase in average investment during the Implementation Stage.

The government's objective is common knowledge to all agents at date 0. It is assumed that the threshold  $b^* \in (0, \bar{b})$ , which implies the government's policy decision is influenced by the econometric evidence. We then consider Rational Expectations Equilibria in which the government, aided by its in-house econometrician, is able to infer  $\mathbf{b}$  based on the econometric evidence.

Let  $\chi$  be an indicator function for average private benefits exceeding the critical threshold for deregulation:

$$\frac{1}{M} \sum_{m=1}^M b_m \geq b^* \Leftrightarrow \chi(\mathbf{b}) = 1. \quad (6)$$

The policy beliefs for firms in industry  $m$  are captured by a function  $\beta$  reflecting their updated assessment of the probability of deregulation during the Implementation Stage, with the updating based upon the realized value of own-industry benefits. We have:

$$\beta(b) \equiv \Pr[\pi_I^*(\mathbf{b}) = \bar{\pi} | b_m = b] = \int_0^{\bar{b}} \dots \int_0^{\bar{b}} \int_0^{\bar{b}} \dots \int_0^{\bar{b}} \left[ \begin{array}{c} [\chi(b_1, \dots, b_{m-1}, b, b_{m+1}, \dots, b_M)] \\ [F(db_1) \dots F(db_{m-1}) F(db_{m+1}) \dots F(db_M)] \end{array} \right]. \quad (7)$$

Notice, under the stated assumptions, beliefs have the same functional form ( $\beta$ ) for all industries. However, the realized value of the function's argument  $b$  will differ across industries with probability one.

The following lemma, which follows directly from equation (7), summarizes some important intuitive properties of the belief function.

**Lemma 1** *If government can implement optimal policies in response to experimental evidence, a firm's assessment of the probability of long-term deregulation is non-decreasing in its own causal parameter  $b_j$ , and strictly increasing on a set of positive measure. Heterogeneous causal parameters ( $\mathbf{b}$ ) result in endogenously heterogeneous policy beliefs on a set of positive measure.*

The intuition for the preceding lemma is simple. A firm's knowledge that its industry faces a high cost ( $b$ ) of regulation rationally assigns a high probability to the government deregulating in the long-term once policymakers have inferred the true realized vector of industry-level cost parameters using the experimental evidence.

To illustrate, consider the simplest setting where the economy has only two industries ( $M = 2$ ). Applying equation (7) we have:

$$\beta(b) = 1 - F(2b^* - b) \Rightarrow \beta'(b) = f(2b^* - b) \geq 0. \quad (8)$$

It is apparent from the preceding equation that, consistent with the preceding lemma, the belief function is non-decreasing. Moreover, different realized values of the causal parameters  $(b_1, b_2)$  tend to result in endogenously heterogeneous policy beliefs 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  $b^*$ .

Continuing with our two-industry example, Figure 1 plots how a firm's assessment of the probability of long-term deregulation will vary with the realization of its own-industry benefit parameter  $b$ . The baseline case assumes  $F$  is the uniform distribution on  $[0, 1]$  and considers a government that will deregulate if average benefits exceed  $b^* = 0.30$ . The figure also considers the effect on the belief function of a more demanding threshold for deregulation,  $b^* = 0.60$ . Finally, the figure considers the effect on beliefs of the higher deregulation threshold in conjunction with a less favorable prior distribution, with  $F$  being a triangular distribution on  $[0, 1]$  with its mode at 0. It is apparent that changes in the government's objective function and changes in priors lead to shifts in the belief function.

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 post-experiment. 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. 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 true value of the causal parameters ( $\mathbf{b}$ ) that one is hoping to estimate.

Although the details will differ, the type of endogenously heterogeneous policy expectations described above are likely to be ubiquitous in settings with evidence-based policymaking: E.g. a measure zero firm suffers a great deal under an experimental regulation, 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 relax the regulation.

## 2.4 Investment Decisions

The model is formally solved in the appendix via optimal control. This subsection confines attention to characterizing firm investment. With analytical expressions for investment in-hand, analytical expressions will be derived for the experimental outcome variable in NPEs, the change in investment

at the start of the experiment, and in RCTs, the difference between the investments of regulated and unregulated firms.

The model solution boils down to three stages, and is thus akin to a three period model. Consider first the Endogenous Policy Economy. At each instant, optimal investment equates the marginal investment cost with the shadow value of a unit of installed capital, call it  $q$ . Therefore, the following function  $i^*$  maps  $q$  at each instant to optimal investment:

$$q = \psi'(i^*) \Rightarrow i^*(q) \equiv \left( \frac{\nu - 1}{\nu\gamma} \right)^{\nu-1} q^{\nu-1}. \quad (9)$$

Optimal investment is increasing in  $q$ , with the  $q$ -sensitivity of investment varying with the parameters of the investment cost function,  $\gamma$  and  $\nu$ .

Since capital depreciates at rate  $\delta$ , the discounted value derived from one unit of capital at its installation date, call this date  $\tau$ , is given by:<sup>4</sup>

$$\begin{aligned} q(x_\tau, b) &= \int_0^\infty e^{-rt} E_\tau \left\{ e^{-\delta t} \Omega_k(k_{\tau+t}, x_{\tau+t}, \pi_{\tau+t}, i_{\tau+t}, b) \right\} dt \\ &= \int_0^\infty e^{-(r+\delta)t} E_\tau \{ x_{\tau+t} + b\pi_{\tau+t} \} dt. \end{aligned} \quad (10)$$

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

Using the Gordon Growth Formula to value the perpetual claim to  $x$ , it follows that the shadow value of capital during each stage  $S \in \{P, E, I\}$  can be expressed as follows:

$$q^S(x, b) = \frac{x}{r + \delta - \mu} + \eta^S(b). \quad (11)$$

In the preceding equation,  $\eta^S(b)$  represents the present value of the flow of private benefits in state  $S$  for a firm with benefit parameter  $b$ . The private benefit value function is derived next.

During the Implementation Stage, private benefits represent a constant perpetuity. Applying the effective discount rate  $r + \delta$ , we have:

$$\eta^I(b) = \frac{\pi_I b}{r + \delta}. \quad (12)$$

Throughout, we let upper (lower) bars denote values and policies if during the current stage  $\pi = \bar{\pi}$  ( $\pi = \underline{\pi}$ ). We thus have the following respective shadow value expressions under deregulation and

---

<sup>4</sup>For a derivation, see the appendix or Abel and Eberly (1994, 1997).

regulation during the Implementation Stage:

$$\begin{aligned}\bar{q}^I(x, b) &= \frac{x}{r + \delta - \mu} + \frac{\bar{\pi}b}{r + \delta} \\ \underline{q}^I(x, b) &= \frac{x}{r + \delta - \mu} + \frac{\underline{\pi}b}{r + \delta}\end{aligned}\tag{13}$$

Consider next the present value of future private 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 private benefits plus the capital gain accruing if there is a transition to the Implementation Stage. Thus, we have the following equilibrium condition:

$$\begin{aligned}(r + \delta)\eta^E(b) &= \pi_E b + \lambda_I [\beta(b)\bar{\eta}^I(b) + (1 - \beta(b))\underline{\eta}^I(b) - \eta^E(b)] \\ \Rightarrow \eta^E(b) &= \left[ \frac{(r + \delta)\pi_E + \lambda_I[\beta(b)\bar{\pi} + (1 - \beta(b))\underline{\pi}]}{(r + \delta)(r + \delta + \lambda_I)} \right] b.\end{aligned}\tag{14}$$

We thus have the following respective shadow value expressions for firms that are deregulated and regulated during the Experiment Stage:

$$\begin{aligned}\bar{q}^E(x, b) &= \frac{x}{r + \delta - \mu} + \left[ \bar{\pi} + \lambda_I \left( \frac{\beta(b)\bar{\pi} + (1 - \beta(b))\underline{\pi}}{r + \delta} \right) \right] \left[ \frac{b}{r + \delta + \lambda_I} \right] \\ \underline{q}^E(x, b) &= \frac{x}{r + \delta - \mu} + \left[ \underline{\pi} + \lambda_I \left( \frac{\beta(b)\bar{\pi} + (1 - \beta(b))\underline{\pi}}{r + \delta} \right) \right] \left[ \frac{b}{r + \delta + \lambda_I} \right].\end{aligned}\tag{15}$$

Consider finally the present value of private benefits evaluated during the Pre-Experiment Stage. The return here consists of the flow of private benefits plus the capital gain accruing if there is a transition to the Experiment Stage. Accounting for the fact that firms face the prospect of deregulated status ( $\pi_E = \bar{\pi}$ ) with probability  $\theta$  during the experiment, we have the following equilibrium condition:

$$\begin{aligned}(r + \delta)\eta^P(b) &= \pi_P b + \lambda_E [\theta\bar{\eta}^E(b) + (1 - \theta)\underline{\eta}^E(b) - \eta^P(b)] \\ \Rightarrow \eta^P(b) &= \left[ \frac{\pi_P}{r + \delta + \lambda_E} + \lambda_E \left( \frac{(r + \delta)(\theta\bar{\pi} + (1 - \theta)\underline{\pi}) + \lambda_I(\beta(b)\bar{\pi} + (1 - \beta(b))\underline{\pi})}{(r + \delta)(r + \delta + \lambda_I)(r + \delta + \lambda_E)} \right) \right] b.\end{aligned}\tag{16}$$

Thus, as a general matter:

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

For example, we have the following respective shadow value expressions under full deregulation ( $\bar{\pi}$ ) and full regulation ( $\underline{\pi}$ ) during the Pre-Experiment Stage:

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

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

$$\beta(b)\bar{\pi} + [1 - \beta(b)]\underline{\pi} \rightarrow \pi_I^{EX}. \quad (18)$$

In the Exogenous Policy Economy, beliefs are homogeneous across firms and industries. And relative to firms in the Endogenous Policy Economy, policy beliefs are extreme in the sense that  $\beta$  is either 0 or 1.

Before closing this subsection it is worth pinning down the causal effect ( $CE$ ) parameters in this economy. As stated by Heckman (2000), the formal definition of a *causal effect* is a Marshallian comparative static. In our economy, the causal effect of interest to the government is how each industry's investment will vary with the regulatory policy put into place during the Implementation Stage. From equation (13) it follows that the causal effect for industry  $m$  is:

$$CE_m \equiv \frac{\partial i^I(x, b_m; \pi)}{\partial \pi} = \left( \frac{\partial q^I(x, b_m; \pi)}{\partial \pi} \right) \left( \frac{di^*}{dq} \right) = \left( \frac{b_m}{r + \delta} \right) \left( \frac{di^*}{dq} \right). \quad (19)$$

From the preceding equation it follows that correctly forecasting firm responses to changes in long-term regulatory policy ( $\pi_I$ ), requires correctly inferring the vector  $\mathbf{b}$  of causal effect parameters. This is the inference problem we analyze below.

### 3 Inference in Natural Policy Experiments

This section considers causal parameter inference in the context of Natural Policy Experiments (NPEs). In an NPE, all firms in the economy face a common exogenous policy shock at a random date.

#### 3.1 Barriers to Inference in NPEs

The econometric challenges in NPEs are best illustrated initially by considering specific examples. We provide analytic characterizations below.

Consider an economy with two industries with the following parameter values:  $r = .05$ ;  $\delta = .10$ ;  $\mu = 0$ ;  $\lambda_E = .15$ ;  $\lambda_I = .15$ ;  $\bar{\pi} = 1$ ;  $\underline{\pi} = 0$ ;  $x = 1$ ;  $\gamma = 1$ ; and  $\nu = 2$ . We consider long expected policy regime durations of  $\lambda^{-1} = 6.7$  years, approximating reality for economy-wide policy changes. Since investment costs are quadratic ( $\nu = 2$ ), optimal investment (equation (9)) is linear in the shadow value of capital  $q$ . The baseline case assumes private benefits ( $b$ ) are drawn from a uniform distribution on  $[0, 1]$ , with the government fully deregulating post-experiment if average private

benefits are greater than its adopted threshold here assumed to be  $b^* = 0.30$ . Beliefs regarding the probability of long-term deregulation for the baseline case were depicted in Figure 1.

We begin first with an example illustrating how ex post policy endogeneity opens up the potential for faulty inference regarding causal effect *signs*. Consider then an NPE involving the experimental imposition of tighter regulation. Specifically, before the experiment, all firms in the economy are partially deregulated, facing  $\pi_P = 0.20$ . When the experiment arrives at an exogenous random date, all firms in the economy will be fully regulated, with  $\pi_E = \underline{\pi}$ . After the experiment, the government will implement its preferred policy.

Given the technology here, one might expect firms to reduce investment when faced with regulatory tightening of the sort taking place during the posited experiment. However, turning to Figure 2A, one sees that firms in this economy will actually *increase* investment in response to the imposition of regulatory tightening during the natural experiment. Intuitively, firms here have a rational expectation of there being a high probability (see Figure 1) of the government finding it optimal to fully deregulate after it learns the true value of the causal parameters using the experimental evidence. Moreover, as shown in Figure 2A, since the rational firm-level belief regarding the probability of long-term deregulation is increasing in  $b$ , the positive investment reaction to the experiment in regulatory tightening is actually increasing in the private benefit parameter  $b$ . Confronted with such empirical evidence, it would perhaps be tempting to search for some causal theory for why firms would benefit from regulation, e.g. regulatory capture. However, the reality in this economy is that regulation is, by construction, harmful to firms, with ex post policy endogeneity clouding the causal effect.

Figure 2B illustrates the potential for bias in inferred causal effect *magnitudes*. Specifically, consider now an NPE featuring a switch from full deregulation ( $\pi_P = \bar{\pi}$ ) prior to the experiment followed by full regulation ( $\pi_E = \underline{\pi}$ ) during the experiment. In the Endogenous Policy Economy, the government's econometrician will correctly infer  $\mathbf{b}$  based on responses to the shock and the government then implements its optimal policy. In the Exogenous Policy Economy, the government is assumed to be powerless to undo the policy shock, so during the Implementation Stage  $\pi_I^{EX} = \underline{\pi}$ .

Consider now inference by an econometrician sitting outside the government, say, an academic economist. The issues are illustrated in Figure 2B. The horizontal axis represents the true private benefit, the causal parameter to be inferred. The vertical axis measures the investment decrease that occurs when the experimental regulation is imposed. It is useful to consider first the contrast between responses in the Endogenous versus Exogenous policy economies (solid versus dashed schedules). In both economies, there is no change in investment if  $b = 0$ . After all, if total profits (equation (3)) are invariant to regulation, investment will also be invariant. Further, in this particular example, firms cut investment by more the higher the flow cost ( $b$ ) of regulation. Monotonicity implies the

econometrician can invert the response function and infer the true causal parameter value, but only if she correctly accounts for policy feedback.

With this in mind, suppose the academic econometrician works in the Endogenous Policy Economy, the sort of economy envisioned by some randomization advocates, where evidence informs policy. If the econometrician is sophisticated, she will anticipate the utilization of econometric evidence as a policy input and read off the solid curve when performing inference, resulting in correct estimation of  $b$ . If she is naïve, policy feedback 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 regulation is permanent with probability 1. It is apparent the naïve econometrician’s parameter estimates will here be biased downward. For example, if the true realized  $b = 1$ , the observed investment change will be equal to  $-.80$ . However, incorrectly reading off the dashed line, the naïve econometrician will infer  $b = .45$ , a downward bias of 55%. Intuitively, in the Endogenous Policy Economy, the government will often find it optimal to reverse Nature’s course and reinstate deregulation based on the experimental data. Consequently, the negative investment response to regulation shock will be less dramatic at each value of  $b$ .

The preceding argument can be recast in heuristic terms. The econometrician who incorrectly reads off the dashed line is working under the premise that the experimental regulation is permanent. If the investment reaction to regulation is zero (close to zero) the bias arising from this incorrect working assumption is zero (small). Suppose instead the econometrician observes a large investment reduction and infers that  $b$  must be high. She might well ask herself, “The firms know that the cost of this regulation is high. They also know that this evidence-based government will eventually figure this out and respond by relaxing the regulation. Well, maybe my working premise that the regulation is permanent, and that firms are acting under this same premise, is not so sensible after all.” All we have done here is formalize this line of reasoning so that the interpretation of the experimental evidence is internally consistent with the role it will play in the policy setting process.

When econometricians examine responses to policy changes, a common concern is that the observed policy change was discretionary, creating the need to closely scrutinize governmental objectives. The hope is that studying exogenous policy shocks allows one to avoid making assumptions about governmental objectives. However, it is apparent that correct causal parameter inference in NPEs still requires a correct stipulation of the government objective function.

To illustrate, suppose the government now demands that in order to deregulate the average firm-level benefit must be still higher than in the baseline scenario, with  $b^*$  raised from  $.30$  to  $.60$ . As shown in the dotted-dashed line in Figure 2, for each value of the causal parameter, the investment contraction will be more severe than in the baseline. Intuitively, with the higher evidence threshold, firms rationally assess a lower probability of deregulation post-experiment, resulting in a more severe

investment contraction. The failure to correctly account for the change in governmental objectives would now result in causal parameter estimates to be biased upwards. For example, if  $b = .60$  the observed investment change will be equal to  $-.80$ . However, incorrectly reading off the original solid curve, the naïve econometrician will infer  $b = 1$ , an upward bias of 67%.

Another common concern expressed by econometricians is that selection limits external validity. For example, it might be thought that a novel regulation is more likely to be imposed by a government that knows its economic environment is unusually benign (low  $\mathbf{b}$ ). Assumption 1 precludes this. Nevertheless, ex post endogeneity leads to similar issues. To illustrate, suppose there is an otherwise equivalent economy in which prior beliefs regarding causal parameters are less favorable to deregulation post-experiment, with  $F$  being a triangular distribution on  $[0, 1]$  with mode 0. This case is captured by the dotted line in Figure 2. Notice, for any given realized value of the causal parameter  $b$ , the investment contraction is more severe. Intuitively, under these new negative prior beliefs, firms assign lower probability to deregulation post-experiment and so respond more aggressively to the experimental regulation. It follows that correct causal parameter inference requires a correct stipulation of prior beliefs regarding the distribution of these same parameters. Moreover, the validity of extrapolating evidence across two economies requires that agents in the two economies hold in common this prior distribution, a strong assumption.

The challenges to inference described above are readily shown analytically. The investment response function ( $R$ ) associated with the introduction of an experimental shift from full deregulation to full regulation is given by:

$$\begin{aligned} R(b) &\equiv i^*[q^E(x, b)] - i^*[\bar{q}^P(x, b)] \\ &= \left(\frac{\nu - 1}{\nu\gamma}\right)^{\nu-1} \left[ (q^E(x, b))^{\nu-1} - (\bar{q}^P(x, b))^{\nu-1} \right] \end{aligned} \tag{20}$$

where the shadow values ( $q$ ) just before and just after the start of the experimental regulation are as shown in equations (15) and (17).

A pre-requisite for parameter inference here is that the regulation response function  $R$  defined above be strictly monotone in  $b$ , which may or not hold, as we later demonstrate. Formally, inference based on an NPE entails inverting the response function  $R$ , allowing one to write:

$$b = R^{-1}[R(b)].$$

When inspecting equation (20), it is important to note that beliefs ( $\beta(b)$ ) enter as arguments into the respective expressions for  $q$ . It follows that any change in economic environment ( $F$  or  $b^*$ ) shifting belief functions shifts the regulation response function  $R$ . As shown in the examples above, failure to account for such shifts results in bias.

To illustrate the issues most clearly, assume investment costs are quadratic ( $\nu = 2$ ). The regulation response function (20) simplifies as follows:

$$R(b) = - \left( \frac{1}{2\gamma} \right) \frac{(r + \delta)(\bar{\pi} - \underline{\pi}) + \lambda_I[\bar{\pi} - \beta(b)\bar{\pi} - (1 - \beta(b))\underline{\pi}]}{(r + \delta + \lambda_I)(r + \delta + \lambda_E)} \quad (21)$$

It is apparent from the preceding equation that a correct stipulation of the belief function  $\beta$  is required to correctly infer the true causal parameter  $b$  based upon a measured responses to the random imposition of regulation. In turn, as shown in Subsection 2.3, a correct stipulation of the belief function requires a correct stipulation of priors ( $F$ ) over the causal parameters to be estimated, as well as a correct stipulation of the government's objective function ( $b^*$ ).

Differentiating the regulation response function in equation (21) one obtains:

$$R'(b) = \left( \frac{1}{2\gamma} \right) \left[ \left( \frac{\lambda_I(\bar{\pi} - \underline{\pi})b}{(r + \delta + \lambda_I)(r + \delta + \lambda_E)} \right) \beta'(b) - \frac{(r + \delta)(\bar{\pi} - \underline{\pi}) + \lambda_I[\bar{\pi} - \beta(b)\bar{\pi} - (1 - \beta(b))\underline{\pi}]}{(r + \delta + \lambda_I)(r + \delta + \lambda_E)} \right]. \quad (22)$$

From the preceding equation and the fact that  $\beta' \geq 0$  (Lemma 1) it follows that endogenously heterogeneous policy beliefs serve to attenuate cross-sectional differences in responses to the imposition of experimental regulation. Intuitively, a firm experiencing a high draw of  $b$  rationally anticipates the government will be more likely to deregulate after learning from the experimental evidence, thus attenuating its negative response to the imposition of experimental regulation. Conversely, the same argument implies that the beliefs channel would amplify the cross-sectional heterogeneity of responses to a deregulation shock.

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

$$\underset{Endogenous}{R(b)} - \underset{Exogenous}{R(b)} = \left( \frac{\lambda_I}{2\gamma} \right) \left[ \frac{\beta(b)\bar{\pi} + (1 - \beta(b))\underline{\pi} - \pi_I^{EX}}{(r + \delta + \lambda_I)(r + \delta + \lambda_E)} \right] b. \quad (23)$$

It follows from equation (23) 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, with  $\pi_I^{EX} \in \{\underline{\pi}, \bar{\pi}\}$ . In the latter economy, post-experiment policy is random, reflecting the uncertain outcome of the econometric inference.

The following proposition summarizes the results of this subsection.

**Proposition 1** *Response functions ( $R$ ) for experimental regulation differ according to whether evidence is relevant (endogenous policy post-experiment) or irrelevant (exogenous policy post-experiment). Across economies with endogenous policies post-experiment, response functions are equal if and only if they share common government objective functions and prior beliefs ( $F$ ) regarding the distribution of causal parameters to be estimated. Endogenous belief heterogeneity decreases (increases) cross-sectional heterogeneity of responses to regulation (deregulation) shocks.*

### 3.2 Observer Effects in NPEs

This subsection offers a simple example illustrating inherent feedback between the perceived credibility of evidence from an NPE and the nature of the evidence itself.

To illustrate, consider two economies facing the same experimental regulation discussed in the preceding subsection. Assume now that the governments in the two economies both have the ability to choose long-term policy,  $\pi_I \in [\underline{\pi}, \bar{\pi}]$ , at their discretion. The only difference between the two economies is that their respective governments have different views on the credibility of experimental evidence. In Economy C, the government views experimental evidence as credible. In Economy NC, the government views experimental evidence as non-credible.

Suppose 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  $\pi_{IP}^*$ . And if no observation occurs, the government has no choice but to rely on its priors, so it again implements  $\pi_{IP}^*$ . 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. Since expectations are unaffected by the act of observation here, no Hawthorne Effect emerges.

In contrast, in Economy C, the distribution of long-term policy varies according to whether or not observation occurs. If no observation occurs, the government is forced to rely on priors, implementing  $\pi_I = \pi_{IP}^*$  just as in Economy NC. However, if observation occurs, the government views the incoming econometric evidence as credible and uses it to infer the parameter vector  $\mathbf{b}$ . Firms then rationally anticipate the government implementing  $\pi_I^*(\mathbf{b})$ , as defined in equation (5), and thus form industry-specific beliefs in accordance with  $\beta$  (equation (7)). 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 here can be expressed in terms of the shadow value of capital. Accounting for the change in beliefs regarding policy during the Implementation Stage, we have:

$$[\underbrace{q^E(x, b)}_{\text{Observed}} - \underbrace{\bar{q}^P(x, b)}_{\text{Not Observed}}] = [\underbrace{q^E(x, b)}_{\text{Not Observed}} - \underbrace{\bar{q}^P(x, b)}_{\text{Not Observed}}] + \underbrace{\frac{\lambda_I[\beta(b)\bar{\pi} + (1 - \beta(b))\underline{\pi} - \pi_{IP}^*]b}{(r + \delta + \lambda_I)(r + \delta + \lambda_E)}}_{\text{Hawthorne Effect}}. \quad (24)$$

We 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 illustrates. The figure assumes private benefits are drawn from a uniform distribution on  $[0, 1]$ , with the government adopting a deregulation threshold  $b^* = 0.45$ . The figure plots experimental responses in an economy where NPEs are viewed as credible sources of evidence. The solid line depicts the regulation response function if firms are not observed and the dashed line depicts the response function if firms are observed. If firms are not observed, the government sets long-term policy based upon prior beliefs and deregulates with probability one ( $\pi_{IP}^* = \bar{\pi}$ ) since the cutoff for deregulation is less than the unconditional average of the private benefit. With observation, sophisticated analysis allows the government to infer  $\mathbf{b}$  and so it implements a contingent optimal policy, with  $\pi_I = \pi_I^*(\mathbf{b})$ . Notice, if firms are observed, they cut their investment by a relatively large amount, especially if their private benefit is low, since in this case they attach a relatively low probability to the regulation being reversed in the long-term.

## 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 considered are stripped of standard self-selection and endogeneity concerns (Assumption 1).

### 4.1 Inference in RCTs

Although analytic characterizations are provided below, the econometric challenges in RCTs are best illustrated initially by considering a specific example.

Consider the following RCT. Prior to the experiment, there is no regulation and all firms face  $\pi_P = \bar{\pi}$ . During the experiment, an equal measure ( $\theta = 1/2$ ) of firms are assigned to the two possible  $\pi$  values. We suppose that 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  $\pi_I = \bar{\pi}$ . In the Endogenous Policy Economy, the government sets policy optimally given the evidence supplied by its econometrician.

For the purpose of numerical illustration, assume two industries, with private benefits being i.i.d. draws from the uniform distribution on  $[0, 1]$ . Government will deregulate if the average private 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:  $r = .05$ ;  $\delta = .10$ ;  $\mu = 0$ ;  $\lambda_E = 1$ ;  $\lambda_I = 1$ ;  $\bar{\pi} = 1$ ;  $\underline{\pi} = 0$ ;  $x = 1$ ;  $\gamma = 1$ ; 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 (unregulated) group ( $\pi_E = \bar{\pi}$ ) and the treatment (regulated) 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, specifically a triangular distribution on  $[0, 1]$  with mode at 1. The dotted line considers the effect of a higher cutoff threshold for deregulation ( $b^* = .85$ ).

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. 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  $b = .70$ .

An 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 the 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 the economy 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 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. 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.

The difference between the observation and non-observation states is due to differences in 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 ( $\pi_E = \bar{\pi}$ ) and treatment ( $\pi_E = \underline{\pi}$ ) groups. Using equation (15) we have:

$$\bar{q}^E(x, b) - \underline{q}^E(x, b) = \frac{(\bar{\pi} - \underline{\pi})b}{(r + \delta + \lambda_I)} \quad (25)$$

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 random 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  $\bar{q}^E$  and  $\underline{q}^E$  shown in equation (25) is just the present value of a claim to the flow of excess private benefits  $(\bar{\pi} - \underline{\pi})b$  accruing to the deregulated group during the Experiment Stage.

But recall, in the preceding subsection (Figure 5), the control-treatment investment difference varied along with 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 (25)). It is the behavior of investment, not shadow values, that the econometrician must understand.

To this end, let  $\Delta$  denote the difference between control and treatment group investment. From equation (15) we have:

$$\begin{aligned}
\Delta(x, b) &= i^*_{Control}[\bar{q}^E(x, b)] - i^*_{Treatment}[q^E(x, b)] & (26) \\
&= \left(\frac{\nu - 1}{\nu\gamma}\right)^{\nu-1} \left[ \left( \underline{q}^E(x, b) + \frac{(\bar{\pi} - \underline{\pi})b}{(r + \delta + \lambda_I)} \right)^{\nu-1} - \left( \underline{q}^E(x, b) \right)^{\nu-1} \right] \\
&= \left(\frac{\nu - 1}{\nu\gamma}\right)^{\nu-1} \left[ \left[ \frac{x}{r + \delta - \mu} + \left( \underline{\pi} + \lambda_I \left( \frac{\beta(b)\bar{\pi} + (1 - \beta(b))\underline{\pi}}{r + \delta} \right) \right) \left( \frac{b}{r + \delta + \lambda_I} \right) + \frac{(\bar{\pi} - \underline{\pi})b}{(r + \delta + \lambda_I)} \right]^{\nu-1} \right. \\
&\quad \left. - \left[ \frac{x}{r + \delta - \mu} + \left( \underline{\pi} + \lambda_I \left( \frac{\beta_H(b)\bar{\pi} + (1 - \beta(b))\underline{\pi}}{r + \delta} \right) \right) \left( \frac{b}{r + \delta + \lambda_I} \right) \right]^{\nu-1} \right].
\end{aligned}$$

A key point to note in equation (26) is that beliefs ( $\beta$ ) 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 (8), 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 (26), 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 one were to assume the investment cost parameter  $\nu$  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 (26) that:

$$\nu = 2 \implies \Delta(x, b) = \left( \frac{1}{2\gamma} \right) \frac{(\bar{\pi} - \pi)}{(r + \delta + \lambda_I)} \times b. \quad (27)$$

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 uncontaminated by 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 policy 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  $\psi$ , where  $\psi$  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-positated 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.<sup>5</sup>

---

<sup>5</sup>See the discussion of Figure 1 in Abel and Eberly (1994).

Recall, under the initially-positied 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 RCT is relevant (endogenous policy post-trial) the treatment group will exhibit a Hawthorne Effect and the control group will exhibit a John Henry Effect. The difference between control and treatment group investment (and the difference in their differences) is invariant to factors affecting post-experiment policy variable expectations if and only if the General Investment Cost Function features: a quadratic adjustment cost function ( $\psi$ ); 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  $\nu$  is equal to 2. The solid line captures both the case of non-observation *and* 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, in this particular case there is no need to account for observation effects when making inferences since the test statistic is unaffected by the act of observation.

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 the earlier NPE which featured an experimental shift from full deregulation to full regulation (Figure 2B). However, we consider now that the government adopts a deregulation threshold of  $b^* = .45$ . Now suppose 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 decrease at the start of the regulation 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 empirical outcome variable 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 the endogeneity of beliefs. Formally, this argument follows from equation (22) which shows that the response to the regulation experiment would be monotone decreasing if the belief function  $\beta$  were constant. However, since  $\beta' \geq 0$ , the response function can be increasing. Intuitively, firms experiencing a high  $b$  value would cut investment more aggressively if their beliefs were the same as other firms. However, such firms rationally attach a lower probability to regulation long-term, which encourages investment. As shown in Figure 7A, this expectations effect can be strong enough to generate a region of  $b$  values, here from 0.80 to 1, where the investment response is actually increasing in  $b$ . Notice, if the observed investment reaction were to fall on the region from 0.60 to 1, the econometrician would not be able to infer the true value of the causal parameter relying exclusively on the observed investment change.

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 the only econometric evidence available to the econometrician 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 econometrician would not be able to invert the empirical outcome variable to solve for the causal parameter.

## 6 Conclusion

This paper illustrates an inherent tension between the credibility of empirical estimates derived from randomization and their practical utilization. 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 policymaker objective function 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 the signs and magnitudes of causal parameters. Far from being stand-alone objects, correct interpretation may require an extremely subtle analysis and may require the imposition of strong functional form assumptions.

The econometric challenges discussed 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 help one 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. Therefore, one may prefer to rely upon other-country evidence or perhaps evidence that is a bit dated, trading off relevance versus contamination. Finally, it would seem to be optimal that experimental subjects do not understand the link between the experiment and subsequent policy decisions. 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 downplay the influence experiments will play in policy decisions. Of course, this is not tenable if the goal is making systematic use of experiments in the policy arena.

## 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] Angrist, Joshua D. and Jorn-Steffen Pischke, 2009, *Mostly Harmless Econometrics: An Empiricist's Companion*, Princeton University Press.
- [5] 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.
- [6] 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.
- [7] 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.
- [8] Dixit, Avinash and Robert Pindyck, 1994, *Investment Under Uncertainty*, Princeton University Press.
- [9] 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.
- [10] 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.
- [11] Hayashi, Fumio, 1982, Tobin's average and marginal q: A neoclassical interpretation. *Econometrica*.

- [12] Heckman, James J., 2000, Causal parameters and policy analysis in economics: A twentieth century retrospective, *Quarterly Journal of Economics* 115 (1), 45-97.
- [13] Hennessy, Christopher A. and Ilya A. Strebulaev, 2015, Beyond random assignment: Credible inference of causal effects in dynamic economies, Working paper, London Business School.
- [14] Keane, Michael P., Structural versus Atheoretic Approaches to Econometrics, *Journal of Econometrics* 156, 1-21.
- [15] 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.
- [16] Ljungqvist, Lars, 2008, Lucas Critique, *The New Palgrave Dictionary of Economics*, Second Edition. Edited by Steven N. Durlauf and Lawrence E. Blume.
- [17] 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.
- [18] Muth, John F., 1961, Rational Expectations and the Theory of Price Movements, *Econometrica* (29), 315-335.
- [19] Rust, John, 2010, Comment on Structural versus Atheoretic Approaches to Econometrics, *Journal of Econometrics* 156, 21-30.
- [20] Sargent, Thomas J., 1971, A note on the accelerationist controversy, *Journal of Money Credit and Banking* 3 (3), 721-725.
- [21] Sargent, Thomas J., 1973, Rational expectations and the dynamics of hyperinflation, *International Economic Review* 14 (2), 328-350.
- [22] Sargent, Thomas J., 1977, The demand for money during hyperinflations under rational expectations, I, *International Economic Review* 18 (1), 59-82.
- [23] Sargent, Thomas J., 1984, Autoregressions, expectations and advice, *American Economic Review* 74, 408-415.
- [24] Sims, Christopher A., 1982, Policy analysis with econometric models, *Brookings Papers on Economic Activity* 1, 107-164.
- [25] Spatt, Chester, 2011, Measurement and policy formulation, working paper, Carnegie Mellon University, Tepper School of Business.

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

## 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  $\nu \in \{2, 3, 4, \dots\}$ . We solve via backward induction.

### *Implementation Stage*

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

$$rV^I(x, k) = \max_i (x + \pi_I b)k - \gamma i^{\nu/(\nu-1)} + \mu x V_x^I(x, k) + \frac{1}{2} \sigma^2 x^2 V_{xx}^I(x, k) + (i - \delta k) V_k^I(x, k). \quad (28)$$

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

$$V^I(x, k) = kq^I(x) + G^I(x). \quad (29)$$

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:

$$(r + \delta)q^I(x) = (x + \pi_I b) + \mu x q_x^I(x) + \frac{1}{2} \sigma^2 x^2 q_{xx}^I(x). \quad (30)$$

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

$$q^I(x) = x\zeta^I + \eta^I. \quad (31)$$

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

$$\begin{aligned} \zeta^I &= \zeta \equiv 1/(r + \delta - \mu). \\ \eta^I &= \frac{\pi_I b}{r + \delta} \end{aligned} \quad (32)$$

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:

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

We begin by noting that

$$q^I(x) = \zeta x + \eta^I \Rightarrow i^*[q^I(x)]q^I(x) - \gamma [i^*(q^I(x))]^{\nu/(\nu-1)} = (\zeta x + \eta^I)^\nu \Gamma = \sum_{h=0}^{\nu} \phi_h^I x^h \quad (33)$$

where

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

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*

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

*has solution*

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

**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:

$$\hat{r}g_h(x) = \mu x g_h'(x) + \frac{1}{2} \sigma^2 x^2 g_h''(x) + \phi_h x^h \quad (34)$$

We conjecture this value function takes the form:

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

Substituting the conjectured solution back into equation (34) one obtains the stated expression for  $\omega_h$ . ■

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

$$\begin{aligned}
G^I(x) &= \sum_{h=0}^{\nu} \phi_h^I \omega_h^I x^h & (35) \\
\text{with } \phi_h^I &\equiv \binom{\nu}{h} (\eta^I)^{\nu-h} \zeta^h \Gamma. \\
\omega_h^I &\equiv \frac{1}{r - \mu h - \frac{1}{2} \sigma^2 h(h-1)}.
\end{aligned}$$

*Experiment Stage*

The HJB equation for the Experiment Stage is:

$$\begin{aligned}
rV^E(x, k) &= \max_i (x + \pi_E b)k - \gamma i^{\nu/(\nu-1)} + \mu x V_x^E(x, k) + \frac{1}{2} \sigma^2 x^2 V_{xx}^E(x, k) & (36) \\
&+ (i - \delta k) V_k^E(x, k) + \lambda_I \beta \left[ \bar{V}^I(x, k) - V^E(x, k) \right] + \lambda_I (1 - \beta) \left[ \underline{V}^I(x, k) - V^E(x, k) \right].
\end{aligned}$$

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

$$V^E(x, k) = kq^E(x) + G^E(x). \quad (37)$$

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:

$$(r + \delta + \lambda_I)q^E(x) = (x + \pi_E b) + \mu x q_x^E(x) + \frac{1}{2} \sigma^2 x^2 q_{xx}^E(x) + \lambda_I \left[ \frac{x}{r + \delta - \mu} + \frac{(\beta \bar{\pi} + (1 - \beta) \underline{\pi})b}{r + \delta} \right]. \quad (38)$$

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

$$q^E(x) = x\zeta^E + \eta^E. \quad (39)$$

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

$$\begin{aligned}
\zeta^E &= \zeta \equiv 1/(r + \delta - \mu). & (40) \\
\eta^E &= \left[ \frac{(r + \delta) \pi_E b + \lambda_I (\beta \bar{\pi} + (1 - \beta) \underline{\pi}) b}{(r + \delta) (r + \delta + \lambda_I)} \right].
\end{aligned}$$

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

$$(r + \lambda_I) G^E(x) = \mu x G_x^E(x) + \frac{1}{2} \sigma^2 x^2 G_{xx}^E(x) + \lambda_I \left[ \beta \bar{G}^I(x) + (1 - \beta) \underline{G}^I(x) \right] + \sum_{h=0}^{\nu} \left[ \binom{\nu}{h} (\eta^E)^{\nu-h} \zeta^h \Gamma \right] x^h. \quad (41)$$

This ODE can be rewritten as:

$$(r + \lambda_I) G^E(x) = \mu x G_x^E(x) + \frac{1}{2} \sigma^2 x^2 G_{xx}^E(x) + \sum_{h=0}^{\nu} \phi_h^E x^h. \quad (42)$$

where

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

From the Growth Option Lemma it follows:

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

*Pre-Experiment Stage*

The HJB equation is:

$$\begin{aligned} rV^P(x, k) &= \max_i (x + \pi_P b)k - \gamma i^{\nu/(\nu-1)} + \mu x V_x^P(x, k) + \frac{1}{2} \sigma^2 x^2 V_{xx}^P(x, k) \\ &+ (i - \delta k) V_k^P(x, k) + \lambda_E \theta \left[ \bar{V}^E(x, k) - V^P(x, k) \right] + \lambda_E (1 - \theta) \left[ \underline{V}^E(x, k) - V^P(x, k) \right]. \end{aligned} \quad (43)$$

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

$$V^P(x, k) = kq^P(x) + G^P(x). \quad (44)$$

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

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

$$(r + \delta + \lambda_E)q^P(x) = (x + \pi_P b) + \mu x q_x^P(x) + \frac{1}{2} \sigma^2 x^2 q_{xx}^P(x) + \lambda_E \theta \bar{q}^E(x) + \lambda_E (1 - \theta) \underline{q}^E(x). \quad (45)$$

We may again conjecture (and verify) the preceding shadow value equation has a linear solution, resulting in equation (17).

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:

$$(r + \lambda_E)G^P(x) = \mu x G_x^P(x) + \frac{1}{2}\sigma^2 x^2 G_{xx}^P(x) + \sum_{h=0}^{\nu} \left[ \binom{\nu}{h} (\eta^P)^{\nu-h} \zeta^h \Gamma \right] x^h + \lambda_E \left[ \theta \overline{G}^E(x) + (1 - \theta) \underline{G}^E(x) \right]. \quad (46)$$

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

$$(r + \lambda_E)G^P(x) = \mu x G_x^P(x) + \frac{1}{2}\sigma^2 x^2 G_{xx}^P(x) + \sum_{h=0}^{\nu} \phi_h^P x^h. \quad (47)$$

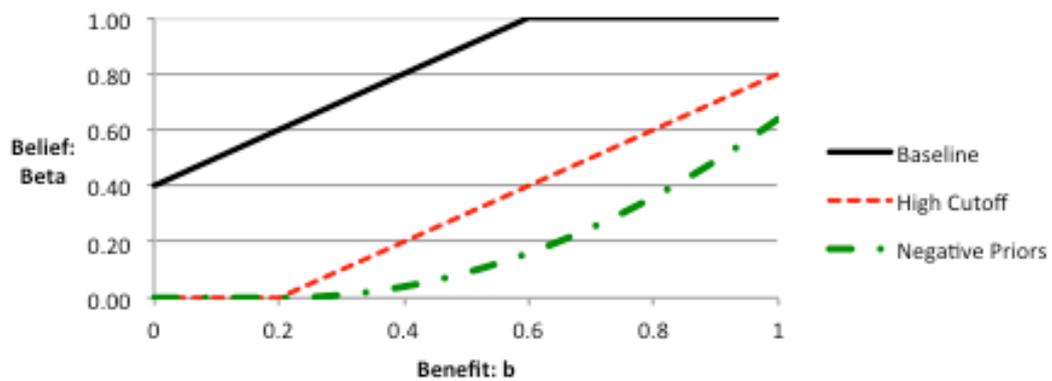
with

$$\begin{aligned} \phi_h^P &\equiv \binom{\nu}{h} (\eta^P)^{\nu-h} \zeta^h \Gamma + \lambda_E \omega_h^E \left[ \theta \overline{\phi}_h^E + (1 - \theta) \underline{\phi}_h^E \right] \\ \overline{\phi}_h^E &\equiv \left[ \lambda_I \left( \beta \overline{\phi}_h^I + (1 - \beta) \underline{\phi}_h^I \right) \omega_h^I \right] + \binom{\nu}{h} (\overline{\eta}^E)^{\nu-h} \zeta^h \Gamma \\ \underline{\phi}_h^E &\equiv \left[ \lambda_I \left( \beta \overline{\phi}_h^I + (1 - \beta) \underline{\phi}_h^I \right) \omega_h^I \right] + \binom{\nu}{h} (\underline{\eta}^E)^{\nu-h} \zeta^h \Gamma. \end{aligned} \quad (48)$$

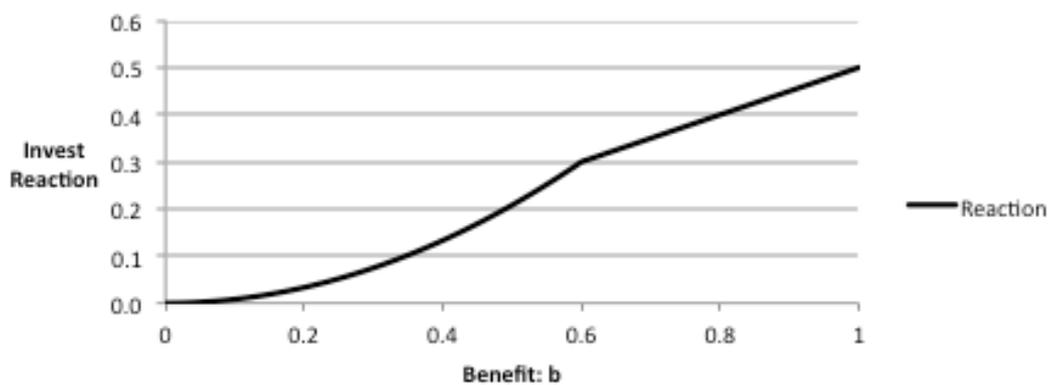
Again, the growth option value is a linear sum of the geometric Brownian motion  $x$  to successive powers. From the Growth Option Lemma it follows that growth option value during the Pre-Experiment Stage is:

$$\begin{aligned} G^P &= \sum_{h=0}^{\nu} \phi_h^P \omega_h^P x^h \\ \omega_h^P &\equiv \frac{1}{r + \lambda_E - \mu h - \frac{1}{2}\sigma^2 h(h-1)}. \end{aligned} \quad (49)$$

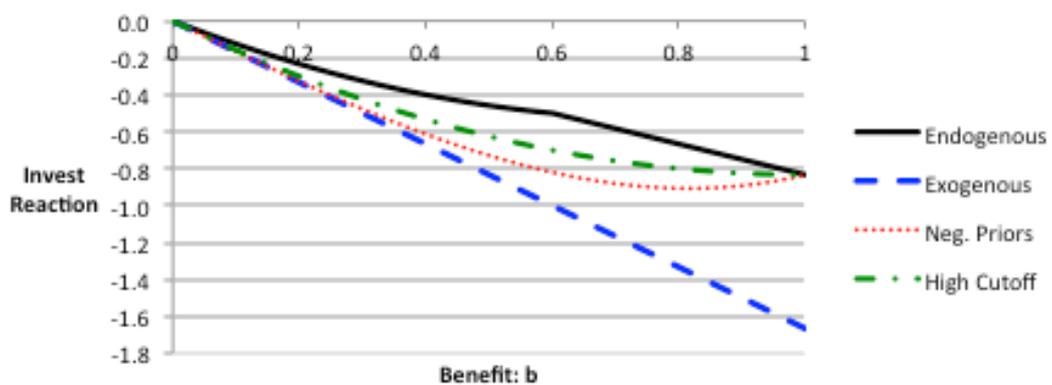
### Figure 1: Endogenous Beliefs



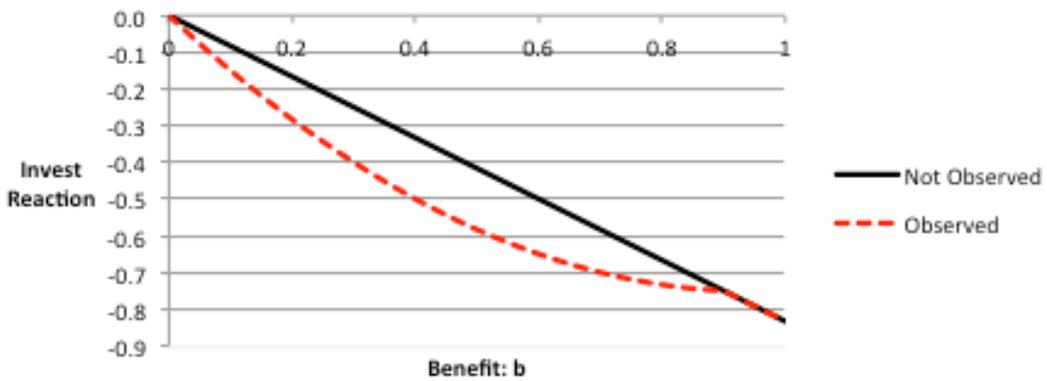
### Figure 2A: Sign Error in Regulation NPE



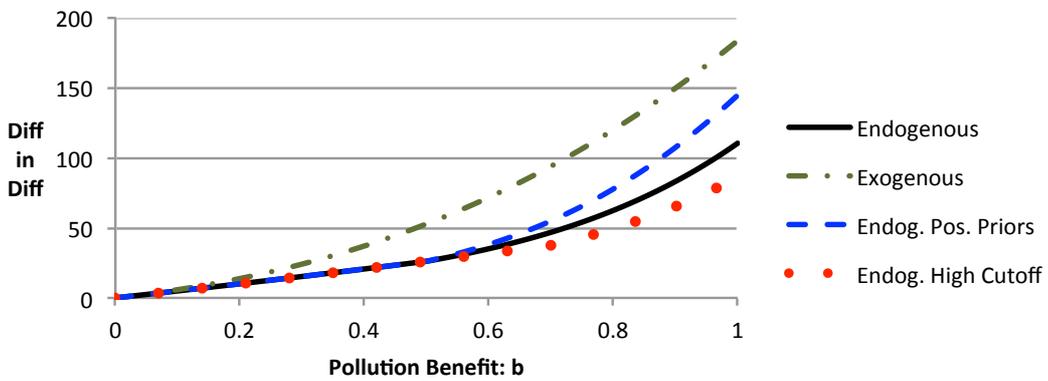
### Figure 2B: Magnitude Bias in Regulation NPE



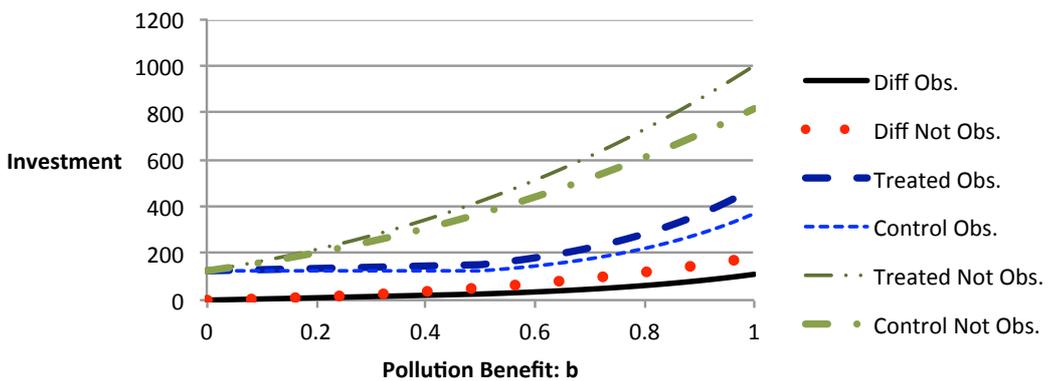
**Figure 3: Hawthorne Effects in NPEs**



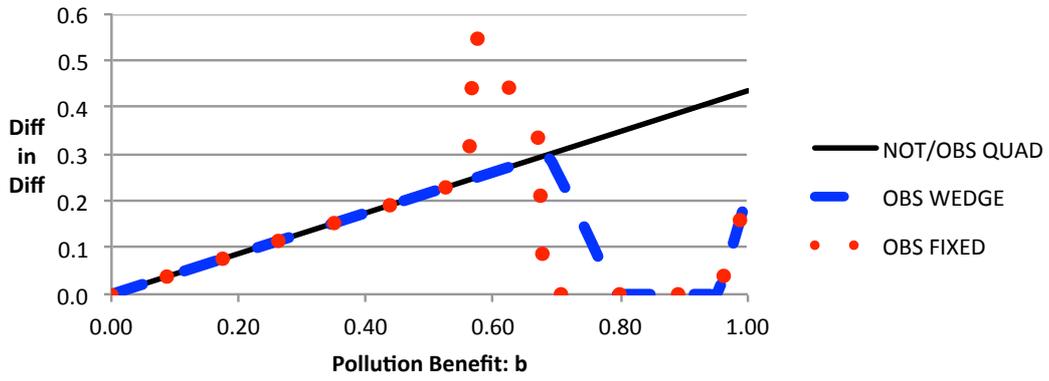
**Figure 4: Randomized Controlled Trial**



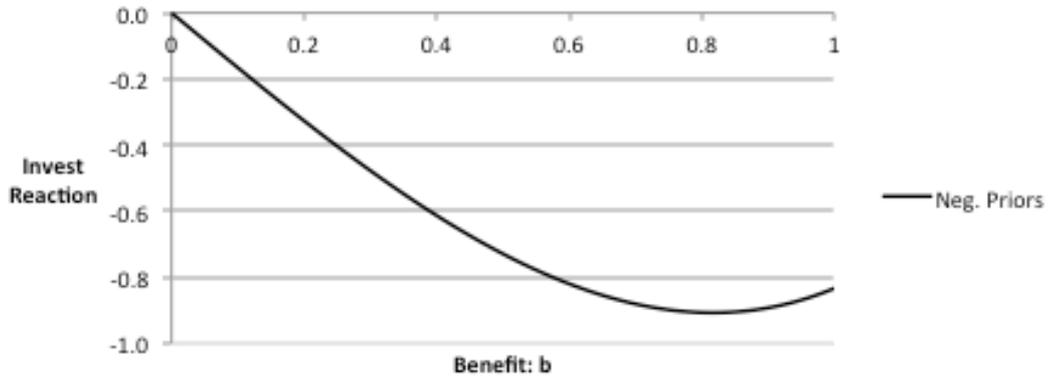
**Figure 5: Observer Effects in RCTs**



### Figure 6: Irreversibility and RCTs



### Figure 7A: NPE Identification Failure



### Figure 7B: RCT Identification Failure

