Beyond Random Assignment: Credible Inference of Causal Effects in Dynamic Economies*

Christopher A. Hennessy
London Business School, CEPR, and ECGI

Ilya A. Strebulaev
Stanford GSB and NBER

September 2015

Abstract

Random assignment is insufficient for measured treatment responses to recover causal effects (comparative statics) in dynamic environments. We provide analytic bias characterizations. If the policy variable is binary (or transitions uniformly distributed) there is attenuation bias. With more than two policy states, treatment responses can undershoot, overshoot, or have incorrect signs. Under permanent random assignment, treatment responses overshoot (have incorrect signs) for realized changes opposite in sign to (small relative to) expected changes. We derive necessary and sufficient conditions, beyond random assignment, for correct inference of causal effects: martingale policy variable or infinitesimal transition rates. Stochastic monotonicity is sufficient for correct sign inference. If these conditions are not met, we show how treatment responses can nevertheless be corrected and mapped to causal effects or extrapolated to forecast responses to future policy changes within or across policy generating processes.

*We thank John Rust, Adriano Rampini and Manuel Adelino (discussant) for constructive input. We also thank seminar participants at Stanford, LBS, Duke, Boston University, LSE, UNC, UBC, NC State, Imperial College, Simon Fraser, Koc, INSEAD, VGSF, SFI, SFS Cavalcade, and the Stanford Conference on Causality in the Social Sciences. Funding from the European Research Council is gratefully acknowledged. This paper was previously circulated under a different title.

1. Introduction

The goal of most empirical work in economics is to estimate signs and magnitudes of causal effects. Heckman (1999) offers the following standard definition of a causal effect.

> Just as the ancient Hebrews were 'the people of the book' economists are 'the people of the model.' ...Within a model, the effects on outcomes of variation in constraints facing agents are well defined. Comparative statics exercises formalize Marshall's notion of a ceteris paribus change which is what economists mean by a causal effect.

Correct empirical estimation of causal effect signs is important given that a theory may be viewed as falsified if it incorrectly predicts a sign. Correct estimation of causal effect magnitudes are important since elasticities are key inputs in welfare calculations. Angrist and Pischke (2010) herald the search for sources of random assignment (or valid instruments) as a "credibility revolution." Their textbook, *Mostly Harmless Econometrics*, states, "The goal of most empirical research is to overcome selection bias, and therefore to have something to say about the causal effect of a variable."

A varied empirical literature exploits (quasi) random treatments in search of causal effects. Greenstone (2002) and Greenstone and Chay (2005) exploit random air quality regulation to estimate causal effects on firm activity and house prices, respectively. Deschenes and Greenstone (2007) use annual weather fluctuations to infer causal effects of long-term climate change on firm profitability. An extensive public finance literature, e.g. Cummins, Hassett and Oliner (1994), treats tax code changes as natural experiments. Romer and Romer (2010) identify exogenous tax changes to estimate their macroeconomic effects. Banerjee, Duflo, Glennerster and Kinnan (2015) and Crepon, Devoto, Duflo and Pariente (2015) estimate causal impacts of microcredit programs on investment using random assignment. Gan (2007) and Chaney, Sraer, Thesmar (2012) exploit real estate price fluctuations to estimate effects of collateral on investment. Werker, Ahmed, and Cohen (2009) use oil price fluctuations to identify causal effects of foreign aid.

The empirical literature often adopts a narrow statistical view regarding estimation. Implicit in much work is the view that "more change is better" since more observations drives down standard errors and permits fixed-effect controls. For example, Slemrod (1992) writes, "Fortunately (for the progress of our knowledge, not for policy), since 1978 the taxation of capital gains has been changed several times, providing much new evidence on the tax responsiveness of realizations." In his analysis of the effect of the Clean Air Act, Greenstone (2001) writes, "The Amendments introduce substantial cross-sectional and longitudinal variation in regulatory intensity at the county level." Cummins, Hassett and Oliner (1994) examine responses to 13 changes in the U.S. corporate income tax from 1962 until 1986. Romer and Romer (2010) identify 54 exogenous tax shocks during the post-war period.

We depart from the extant empirical literature by analyzing the economic meaning of estimates derived from ideal random assignment settings, developing a tractable model laboratory tailor-made for this purpose. The particular theory used as the underlying benchmark is the canonical neoclassical q-theory of investment, e.g. Abel and Eberly (1997). With this benchmark in-hand, we examine the relationship between *causal effects* (theory-implied comparative statics) and measured *treatment responses* (observed responses to government policy changes).

Our analysis shows that in a broad and important class of economic environments, random assignment is insufficient for valid inference. Specifically, in dynamic environments in which policy variables evolve over time stochastically, treatment responses do not generally equal the causal effects empiricists seek to estimate. Notwithstanding the quotations above, and the omission of discussion of biases arising from transience in perhaps the majority of published work and textbooks (e.g. Angrist and Pischke (2009, 2010)), undoubtedly many empiricists have a heuristic sense that policy transience clouds the interpretation of econometric evidence. However, the exact nature of the biases is apparently not well understood. For example, conventional wisdom holds that policy transience results in attenuation bias in the context of real investment. For example, Slemrod (1992), Aaron (1992) and Cummins, Hassett and Oliner (1994) argue policy transience might have led to a downward bias in the estimation of investment elasticities based upon evidence from tax reforms. Atanasov and Black (2015) argue that attenuation bias is a general phenomenon in shock-based inference.

We show that such conventional wisdom is valid if the studied policy variable is binary, the focus of early theoretical work on policy transience. Further, we show that attenuation bias emerges if policy variable transitions are uniformly distributed. Perhaps less well understood is the magnitude of such attenuation bias. For example, under plausible parameterizations, attenuation bias exceeds 50% even if the expected duration of policy regimes is relatively long (10 years).

As we show, the binary policy variable setting is relatively benign. In particular, if the policy variable can take on more than two values, treatment responses can understate, overstate, and even have signs opposite to causal effects. The model allows us to characterize analytically the probability of these alternative bias forms depending on the distribution of randomized treatments, discount rates, and expected policy regime durations. The results are not comforting. For example, consider an economy featuring a 5% discount rate, an expected current regime life of five years, with a future permanent policy variable change drawn from the uniform distribution on [0,1]. Here, the econometrician will face attenuation bias for 60% of the possible random treatments, those on [0.40,1]. For the remaining 40% of possible treatments, those on [0,0.40), the treatment response is opposite in sign to the causal effect. If instead the treatments are drawn from the uniform distribution on [-0.20, 0.80], the treatment response will overshoot the causal effect with probability 20%, have the wrong sign with probability 24%, and undershoot with probability 56%. It is hard to understand the sense in which estimators with such properties can be viewed as credible.

These findings also cast doubt on the interpretation and utilization of elasticity estimates shaping policy. As just one example, in surveying evidence following the Tax Reform Act of 1986, Slemrod (1990) concludes "the short-term response has in most cases been less dramatic than many economists had expected." Slemrod (1992) describes a consensus view of a "downward revaluation of the responsiveness to taxation of real variables." Based on this evidence, Slemrod (1992) and Aaron (1992) call for increased focus on the distributional consequences of taxation, and less focus on excess burdens.

After diagnosing the problems arising from policy transience, the majority of the paper is devoted to a constructive formal analysis of how each of the problems identified can be overcome. We first derive auxiliary assumptions, beyond random assignment, needed to ensure equality of treatment responses and causal effects in dynamic settings. We begin by showing treatment responses converge to causal effects if all policy transition rates tend to zero. However, there are four problems with relying on rare-event transitions. First, since policy changes are outcomes of political processes, they are often anticipated. Second, reliance on rare events implies expected waiting times for relevant natural experiments are extremely long. Third, it is unclear why policymakers should care about changes in a policy variable that almost-surely will not change in the future. Finally, the majority of the empirical literature does not exploit truly rare events.

In actuality, it is not necessary for empiricists to confine attention to rare events. As we show, treatment responses equal causal effects if and only if expected changes in the policy variable have mean zero. Formally, each policy state must be absorbing or, if transitions out of specified states are possible, conditional mean changes must equal zero, i.e. the policy variable must be a martingale. Intuitively, comparative statics represent responses to hypothetical policy changes that are completely unanticipated and permanent. If the policy variable is a martingale, firms extrapolate the current policy state into perpetuity, despite transience. Consequently, they invest as if the current policy state will last forever. Such beliefs and behavior allow empiricists to directly recover causal effects from measured treatment responses even if treatments are transitory.

In certain cases, empiricists are willing to settle for the more limited objective of correct inference of causal effect signs. Here too we derive a number safe-harbor assumptions that empiricists can invoke: binary assignment; uniformly distributed policy transitions; stochastic monotonicity (e.g. Kolokoltsov (2011)) of the policy variable; or permanent random assignment opposite in sign to the expected assignment.

We also show how the outlined problems arising from policy transience can be overcome when the auxiliary assumptions described above are not satisfied. That is, we show how causal effects can be recovered from treatment responses for arbitrary policy transition rates. We also show how treatment responses can be extrapolated across different transitions under the same policy generating process, or across different policy processes, overcoming barriers to valid extrapolation. The proposed correction works as follows. With policy transience, the correct measure of "dosage" is not the change in the policy variable itself but the associated change in shadow value. Measured responses to policy variable transitions can be normalized by shadow value changes to capture response-per-dose of shadow value. With this estimate in-hand, responses to prospective shadow value dosages can be computed.

Our paper draws inspiration from Lucas (1976). Lucas does not analyze the relationship between natural experiments and causal effects, nor does he discuss or compute the probabilities of overshooting and sign reversals. Further, Lucas does not derive conditions under which treatment responses recover causal effects or their signs. As in Lucas (1976), we show the limits of external validity of empirical evidence, here of the experimental variety. Moving beyond his critique in relation to forecasting effects of future policy changes, we show there is no a priori reason to think elasticities can be extrapolated *within* the same policy generating process. Constructively, we show how to extrapolate observed treatment responses between and within policy processes, as well as back to causal effects.

Abel (1983) analyzes effects of permanent versus temporary tax policies under perfect foresight. One can view our work as extending his to capture uncertainty and more than two policy states. Auerbach (1986) and Auerbach and Hines (1988) present investment Euler equations under stochastic tax rates. Hassett and Metcalf (1999) present a real options model with a one-time investment which they use to assess whether uncertainty regarding tax credits encourages investment. Gourio and Miao (2008) numerically compare effects of permanent and temporary dividend tax cuts.

Our work is in the spirit of a paper by Chetty (2012) who writes, "The identification of structural parameters of stylized models is one of the central tasks of applied economics." We use a structural model to understand and correct empirical estimates derived from random assignment in dynamic settings, with recovery of adjustment cost parameters being an interim step. Chetty analyzes how to recover structural elasticity parameters if there are transaction costs or inattention leading to inaction regions. Cummins, Hassett and Oliner (1994) estimate adjustment cost parameters based on tax experiments. However, they assume each tax change is a complete surprise and viewed as permanent, despite the fact that taxes changed every other year during their sample period. We show that under rational expectations, this imputation procedure leads to incorrect measurement of policy dosages and biased estimates of elasticities and adjustment costs.

The rest of the paper is as follows. Section 2 develops a theory of dynamic investment, deriving causal effects. Section 3 describes our model laboratory. Section 4 evaluates potential for biases in treatment response estimators. Section 5 derives auxiliary assumptions ensuring equality of treatment responses and causal effects, and also derives bias corrections for environments where these assumptions are not satisfied. Section 6 illustrates limits on both external and internal (within policy process) validity before deriving corrections facilitating extrapolation. Session 7 presents a numerical example.

2. Neoclassical Theory of Capital Demand

In order to set a benchmark we must first derive causal effects of government policies based on some underlying theory. To this end, this section articulates a neoclassical q-theory of capital demand for firms facing taxation and regulation, following closely the model of Abel and Eberly (1997).

2.1. Technology

Agents are risk-neutral and discount cash payments at rate $r > 0$. Time is continuous and uncertainty is modeled by a complete probability space $(\Omega, \mathcal{F}, \mathfrak{M}, \mathbb{F})$ where $\mathfrak{M}$ is the measure and $\mathbb{F} = \{\mathcal{F}_t : t \geq 0\}$ satisfies the standard conditions.

A price-taking firm produces an output flow each instant employing the production function $\zeta n^{\alpha} s^{1-\alpha}$, with the factor share $0 < \alpha < 1$. The productivity process ζ is positive and stochastic. The production input s represents physical capital which can be adjusted at a cost. Capital depreciates at rate $\delta \geq 0$. The production input n is costlessly adjustable and has unit price $p > 0$.

Government intervention in the economy takes two forms. First, the government imposes a linear tax at rate $\tau \geq 0$ on gross operating profits.[1] Second, the government alters p through its intervention in input markets. For example, minimum wage laws affect labor input prices. Environmental laws affect fuel input prices. Tariffs and state-level taxes are often imposed on production inputs. To fix ideas, throughout we consider an econometrician interested in estimating the causal effects of these commonly-studied government interventions.

The firm sells its output at a positive stochastic price ρ. Let $y \equiv \rho\zeta$ and let z be a standard Brownian motion defined on $(\Omega, \mathcal{F}, \mathfrak{M}, \mathbb{F})$. The process y evolves as a geometric Brownian motion:

$$dy = mydt + \nu ydz. \tag{1}$$

Gross operating profits, net of tax, are denoted π. We have:

$$\pi \equiv \max_{n} \quad [yn^{\alpha}s^{1-\alpha} - pn](1-\tau). \tag{2}$$

It follows that utilization of the flexible input is increasing in y and decreasing in p. Specifically, the optimal flexible input is:

$$n^* = \left(\frac{\alpha y}{p}\right)^{\frac{1}{1-\alpha}} s. \tag{3}$$

Although the focus of our analysis is real investment, throughout the paper effects of government policy on employment follow equation (3).

[1] Other tax bases, e.g. income tax with interest deductions, can be modeled but with considerable added complexity such as the need to model leverage policy. See Hennessy, Kashara and Strebulaev (2015).

Gross operating profits, net of tax, can be expressed as:

$$\pi(s, x, p, \tau) = xs\kappa(p, \tau), \tag{4}$$

with

$$\begin{aligned}\kappa(p, \tau) &\equiv (1-\tau)(1-\alpha)\alpha^{\alpha/(1-\alpha)}p^{-\alpha/(1-\alpha)} \\ x &\equiv y^{\frac{1}{1-\alpha}}.\end{aligned} \tag{5}$$

Equation (5) is central to the analysis, showing how the government in our parable economy influences profitability, with κ decreasing in p and τ. Rather than analyze the effects of each policy lever separately, for the remainder of the paper we treat the government as determining κ directly.

From Ito's lemma it follows that x evolves as a geometric Brownian motion, with

$$dx = \mu x dt + \sigma x dz \tag{6}$$

where

$$\begin{aligned}\mu &\equiv \frac{m}{1-\alpha} + \frac{1}{2}\frac{\alpha\nu^2}{(1-\alpha)^2} \\ \sigma &\equiv \frac{\nu}{1-\alpha}.\end{aligned} \tag{7}$$

To ensure bounded valuations, it is assumed $r > \mu$.

The capital stock evolves according to:

$$\begin{aligned}ds &= (a - \delta s)dt \\ s_0 &> 0.\end{aligned} \tag{8}$$

Following Abel and Eberly (1997), it is assumed s_0 is sufficiently large so that full depletion of the stock can be ignored. Each unit of the stock variable can be purchased and sold at a constant price $\psi \geq 0$. We follow Abel and Eberly in assuming the firm faces quadratic costs γa^2, where $\gamma > 0$.

2.2. Optimal Accumulation

The firm can be understood as choosing its instantaneous accumulation policy (a) to maximize the sum of cash flows and expected capital gains. Applying Ito's lemma, the Bellman equation is:

$$rV(s, x) = \max_a \quad (a - \delta s)V_s(s, x) + \mu x V_x(s, x) + \frac{1}{2}\sigma^2 x^2 V_{xx}(s, x) + \kappa xs - \psi a - \gamma a^2. \tag{9}$$

The optimal investment policy solves:

$$\max_{a} \quad aV_s(s,x) - \psi a - \gamma a^2. \tag{10}$$

We conjecture and then verify the value function is of the form:

$$V(s,x) = sq(x) + G(x). \tag{11}$$

The term $sq(x)$ measures the value of installed capital, while G measures growth option value.

Under the conjectured value function, $V_s(s,x) = q(x)$, and the optimal accumulation policy is:

$$a(x) = \frac{q(x) - \psi}{2\gamma}. \tag{12}$$

The instantaneous net gain attributable to accumulation is:

$$\begin{aligned} aq - \psi a - \gamma a^2 &= (q-\psi)^2 \Gamma \\ \Gamma &\equiv \frac{1}{4\gamma}. \end{aligned} \tag{13}$$

Substituting into the Bellman equation the conjectured value function, and the instantaneous gain under optimal accumulation from equation (13), we obtain:

$$\begin{aligned} rsq(x) + rG(x) &= -\delta sq(x) + \mu x s q_x(x) + \mu x G_x(x) + \frac{1}{2}\sigma^2 x^2 s q_{xx}(x) + \frac{1}{2}\sigma^2 x^2 G_{xx}(x) \\ &\quad + \kappa x s + [q(x) - \psi]^2 \Gamma. \end{aligned} \tag{14}$$

Since the Bellman equation must hold point-wise on the state space, the terms scaled by s in the preceding equation must equate. This implies the following ordinary differential equation describing the evolution of the shadow value of capital:

$$(r+\delta)q(x) = \mu x q_x(x) + \frac{1}{2}\sigma^2 x^2 q_{xx}(x) + \kappa x. \tag{15}$$

From the preceding equation and the Feynman-Kac formula it follows:

$$q(x_0) = \mathbb{E}\left[\int_0^\infty e^{-(r+\delta)t} \kappa x_t dt \mid \mathfrak{F}_0\right]. \tag{16}$$

That is, q is the discounted value of the marginal product of capital.

As in Abel and Eberly (1997), we can rule out bubbles causing unbounded valuations as x goes to zero or infinity. We obtain the following solution to equation (15):

$$q^{**}(x) = \frac{\kappa x}{r + \delta - \mu}. \tag{17}$$

The optimal accumulation policy is:

$$a^{**}(x, \kappa) = \frac{1}{2\gamma}\left[\frac{\kappa x}{r + \delta - \mu} - \psi\right]. \tag{18}$$

Note, in the preceding two equations, and throughout the remainder of the paper, double-stars are used to denote variables obtained under the present section's neoclassical theory of investment.

It follows that:

$$\frac{\partial}{\partial \kappa} a^{**}(x, \kappa) = \frac{x}{2\gamma(r + \delta - \mu)}. \tag{19}$$

Note, given the assumed quadratic adjustment costs, causal effects are linear in the government policy lever κ.

A complete model solution requires computing the growth option value function (G). However, since our objective is to analyze causal effects, the policy function (18) is sufficient. The growth option value is derived in the appendix.

3. The Model Laboratory

One would like to determine the conditions under which natural experiments correctly recover the sign and magnitude of theory-implied causal effects. The remainder of the paper treats the model itself as a laboratory. The laboratory is ideal in two respects. First, in contrast to the real-world environment confronting empiricists, we here know the theory-implied causal effect, as shown in equation (19). Second, we consider that the econometrician inhabits an economy where government policy evolves as an independent stochastic process. This independence assumption allows us to abstract from the type of endogeneity bias that is the primary focus of contemporary empirical work.

3.1. The Policy Generating Process

The government policy variable κ variable takes the form of an independent N-state continuous-time Markov chain.[2] Government policy in regime i is κ_i with $i = 1, ..., N$, and $N \geq 2$. The following

[2]See Ross (1996) for an exposition.

indexing convention is adopted:

$$\kappa_1 > ... > \kappa_N.$$

The parameter $\lambda_{ij} \geq 0$ denotes the *transition rate* from regime i to regime j. Let:

$$\Lambda_i \equiv \sum_{j \neq 1} \lambda_{ij}.$$

The amount of time policy remains in regime i before transitioning is denoted T_i. The random variable T_i is exponentially distributed with parameter Λ_i. It follows that the expected duration of regime i is Λ_i^{-1}. If $\lambda_{ij} = 0$ for all j , then $\Lambda_i = 0$ and state i is *absorbing.* If Λ_i tends to ∞, then state i is an *instantaneous state*, using the terminology of Ross (1996). Finally, for a non-absorbing state i, the conditional probability of transitioning into state j, given a transition out of i, is given by:

$$P_{ij} \equiv \frac{\lambda_{ij}}{\sum_{j \neq i} \lambda_{ij}}.$$

3.2. Model Solution

We turn next to a characterization of optimal accumulation. The value of the firm in policy state i is denoted V^i. Accounting for regime changes, and concomitant capital gains, we have the following system of N Bellman equations:

$$\begin{aligned} rV^1(s,x) &= \max_a \quad (a - \delta s)V_s^1(s,x) + \mu x V_x^1(s,x) + \frac{1}{2}\sigma^2 x^2 V_{xx}^1(s,x) \\ &\quad + \sum_{j \neq 1} \lambda_{1j}[V^j(s,x) - V^1(s,x)] + \kappa_1 xs - \psi a - \gamma a^2 \\ &\quad ... \\ rV^N(s,x) &= \max_a \quad (a - \delta s)V_s^N(s,x) + \mu x V_x^N(s,x) + \frac{1}{2}\sigma^2 x^2 V_{xx}^N(s,x) \\ &\quad + \sum_{j \neq N} \lambda_{Nj}[V^j(s,x) - V^N(s,x)] + \kappa_N xs - \psi a - \gamma a^2. \end{aligned} \tag{20}$$

The optimal state-contingent accumulation policy solves:

$$a^i(x) \in \arg\max_a \quad aV_s^i(s,x) - \psi a - \gamma a^2; \; i = 1, ..., N. \tag{21}$$

We conjecture the solution to the Bellman system (20) has the following functional form:

$$V^1(s,x) = q^1(x)s + G^1(x); ...; V^N(s,x) = q^N(x)s + G^N(x). \tag{22}$$

Under the value function conjectured in equation (22), $V_s^i(s,x) = q^i(x)$, and the optimal state-contingent accumulation policy is:

$$a^i(x) = \frac{q^i(x) - \psi}{2\gamma};\ i = 1, ..., N. \tag{23}$$

From equation (23) it follows accumulation will exhibit a jump each time there is a policy regime transition. Such jumps represent the treatment response measured by the econometrician.

Evaluated at the optimal policy, the instantaneous net gain attributable to accumulation is:

$$a^i(x)q^i(x) - \psi a^i(x) - \gamma[a^i(x)]^2 = [q^i(x) - \psi]^2\Gamma. \tag{24}$$

Substituting the accumulation gain from equation (24) and the conjectured value functions into the original system of Bellman equations, we can rewrite the Bellman system as:

$$\begin{aligned}
& \left(r + \delta + \sum_{j\neq 1}\lambda_{1j}\right) q^1(x)s + \left(r + \sum_{j\neq 1}\lambda_{1j}\right) G^1(x) \qquad (25)\\
= \;& \mu x[sq_x^1(x) + G_x^1(x)] + \frac{1}{2}\sigma^2x^2[sq_{xx}^1(x) + G_{xx}^1(x)] \\
& + \sum_{j\neq 1}\lambda_{1j}[q^j(x)s + G^j(x)] + \kappa_1 xs + [q^1(x) - \psi]^2\Gamma \\
& \ldots \\
& \left(r + \delta + \sum_{j\neq N}\lambda_{Nj}\right) q^N(x)s + \left(r + \sum_{j\neq N}\lambda_{Nj}\right) G^N(x) \\
= \;& \mu x[sq_x^N(x) + G_x^N(x)] + \frac{1}{2}\sigma^2x^2[sq_{xx}^N(x) + G_{xx}^N(x)] \\
& + \sum_{j\neq N}\lambda_{Nj}[q^j(x)s + G^j(x)] + \kappa_N xs + [q^N(x) - \psi]^2\Gamma.
\end{aligned}$$

Since the Bellman equations must be satisfied at all points in the state space, the terms scaled by s in each of the preceding equations must equate. Thus, the following system of N equations must be satisfied:

$$\begin{aligned}
\left(r + \delta + \sum_{j\neq 1}\lambda_{1j}\right) q^1(x) &= \mu x q_x^1(x) + \frac{1}{2}\sigma^2x^2 q_{xx}^1(x) + \sum_{j\neq 1}\lambda_{1j}q^j(x) + \kappa_1 x \qquad (26)\\
& \ldots \\
\left(r + \delta + \sum_{j\neq N}\lambda_{Nj}\right) q^N(x) &= \mu x q_x^N(x) + \frac{1}{2}\sigma^2x^2 q_{xx}^N(x) + \sum_{j\neq N}\lambda_{Nj}q^j(x) + \kappa_N x.
\end{aligned}$$

Applying the Feynman-Kac formula to an arbitrary differential equation in the preceding system, it follows:

$$\begin{aligned} q^i(x_0) &= \mathbb{E}\left[\int_0^\infty e^{-(r+\delta+\Lambda_i)t}\left(\kappa_i x_t + \Lambda_i \sum_{j\neq i} P_{ij} q^j(x_t)\right) dt \mid \mathfrak{F}_0\right] \qquad (27) \\ &= \mathbb{E}\left[\int_0^\infty \left(\int_0^T e^{-(r+\delta)t}\kappa_i x_t dt + e^{-(r+\delta)T}\sum_{j\neq i} P_{ij} q^j(x_T)\right)(\Lambda_i e^{-\Lambda_i T})dT \mid \mathfrak{F}_0\right]. \end{aligned}$$

The second expression for the shadow value offered above follows from the first via integration by parts. It states that the shadow value is the expectation over the current policy regime life, which is exponentially distributed, of the net marginal product of capital up to the regime change plus the expectation of the shadow value of capital post regime change.

We conjecture the following solution to system (26):

$$q^i(x) = xc_i;\ i = 1, ..., N. \qquad (28)$$

Substituting the conjectured linear solutions into the system of equations (26), it follows that the shadow value of capital is:

$$\begin{bmatrix} q^1(x) \\ ... \\ q^N(x) \end{bmatrix} = x\underbrace{[\mathbf{T}(r+\delta-\mu)]^{-1}\begin{bmatrix} \kappa_1 \\ ... \\ \kappa_N \end{bmatrix}}_{\equiv \mathbf{c}}. \qquad (29)$$

where $\mathbf{T}(R)$ denotes the following *augmented transition matrix*:

$$\mathbf{T}(R) \equiv \begin{bmatrix} R+\sum_{j\neq 1}\lambda_{1j} & -\lambda_{12} & ... & -\lambda_{1N} \\ -\lambda_{21} & R+\sum_{j\neq 2}\lambda_{2j} & ... & -\lambda_{2N} \\ ... & ... & ... & ... \\ -\lambda_{N1} & -\lambda_{N2} & ... & R+\sum_{j\neq N}\lambda_{Nj} \end{bmatrix} \qquad (30)$$

If there are only two possible policy states, the preceding shadow value expressions can be written as:

$$q^i(x) = \left(\frac{\kappa_i}{r+\delta-\mu} + \frac{\lambda_{ij}(\kappa_j-\kappa_i)}{(r+\delta-\mu)(r+\delta-\mu+\lambda_{ij}+\lambda_{ji})}\right)x. \qquad (31)$$

To complete the model solution, we must also compute the state-contingent growth option value (G^i). However, since our objective is to analyze the relationship between measured treatment responses and causal effects, the policy function in equation (23) and the shadow value vector (29) are sufficient. Derivation of the growth option value is provided in the appendix.

4. Treatment Responses and Causal Effects

This section analyzes the relationship between measured treatment responses and theory-implied causal effects.

From equation (23) it follows that the optimal accumulation policy is

$$\begin{bmatrix} a^1(x) \\ ... \\ a^N(x) \end{bmatrix} = \frac{1}{2\gamma}\left[\begin{bmatrix} q^1(x) \\ ... \\ q^N(x) \end{bmatrix} - \begin{bmatrix} \psi \\ ... \\ \psi \end{bmatrix}\right] \tag{32}$$

with equation (29) determining the state-contingent shadow value of capital.

Consider now the treatment responses that will be observed by the econometrician. The treatment response associated with a transition from state i to j is denoted TR_{ij}. We have:

$$TR_{ij}(x) \equiv a^j(x) - a^i(x) = \frac{1}{2\gamma}[q^j(x) - q^i(x)]. \tag{33}$$

If follows from the preceding equation that transitions across policy states result in jumps in accumulation proportional to the jumps in the shadow value, with shadow value jumps determined by equation (29).

Consider next causal effects. Section 2 showed causal effects are linear. Therefore, the causal effects associated with discrete policy jumps are correctly computed by multiplying the infinitesimal causal effect (equation (19)) by the discrete policy variable change. We have:

$$CE_{ij} \equiv \frac{x}{2\gamma}\left[\frac{1}{r+\delta-\mu}\right](\kappa_j - \kappa_i). \tag{34}$$

4.1. Attenuation Bias

Typically empirical work focuses on biases arising from policy endogeneity, eschewing problems arising from policy transience. However, some empiricists do offer informal discussions of transience-induced biases, citing attenuation bias as a potential concern, e.g. Slemrod (1992), Aaron (1992), Cummins, Hassett and Oliner (1994), and Atanasov and Black (2015).

This subsection considers the problem of attenuation bias more formally. We begin by considering the simplest possible setting, one in which the policy variable can only take on two possible values. The following proposition follows directly from the formula for the shadow value under binary assignment given in equation (31), with treatment responses equal to $\Delta q/2\gamma$.

Proposition 1 *If there are only two possible policy states, the treatment response is*

$$TR_{ij}(x) = \left(1 - \underbrace{\frac{\lambda_{ij} + \lambda_{ji}}{r + \delta - \mu + \lambda_{ij} + \lambda_{ji}}}_{Attenuation}\right) \underbrace{\left(\frac{1}{2\gamma}\right)\left(\frac{x}{r + \delta - \mu}\right)(\kappa_j - \kappa_i)}_{CE_{ij}}.$$

With binary assignment there is always attenuation bias. Further, the bias is quite severe even in settings with long expected policy regime durations. To illustrate, suppose the econometrician inhabits an economy with long expected policy durations, say 10 years ($\lambda_{ij} = \lambda_{ji} = .10$). In this case, attenuation bias exceeds one-half of causal effects under reasonable parameterizations featuring $r + \delta - \mu < .20$. Similarly, the proposition shows that the size of the attenuation bias exceeds one-half under such parameterizations if one considers, say, a permanent transition ($\lambda_{ji} = 0$) out of a policy regime with an expected duration of 5 years ($\lambda_{ij} = .20$) or an unexpected transition (λ_{ij} infinitesimal) to a new regime with an expected 5 year duration ($\lambda_{ji} = .20$). Although Slemrod (1990) is not clear regarding the basis for his priors regarding the expected size of responses to TRA86, such back-of-the-envelope calculations suggest that small responses to the legislation might well have been expected.

Two other implications of Proposition 1 are worth noting. First, it is apparent that if either regime is instantaneous, the treatment response will be infinitesimal and the attenuation bias is roughly 100% of the causal effect. Conversely, if both transition rates are infinitesimal, treatment responses approximate causal effects.

Consider next a setting in which all transition rates are equal, with $\lambda_{ij} = \lambda > 0$ for all $i \neq j$. In this case, each policy regime has an expected life equal to $1/(N-1)\lambda$. Here, if a transition out of an arbitrary state i occurs, the conditional probability of transitioning to each of the remaining states (P_{ij}) is uniform, equal to $1/(N-1)$. As shown in the appendix, the following proposition follows from equation (29) and the fact that $\Delta a = \Delta q/2\gamma$.

Proposition 2 *If all transition rates are equal to λ, with N policy states the treatment response is*

$$TR_{ij}(x) = \left(1 - \underbrace{\frac{N\lambda}{r+\delta-\mu+N\lambda}}_{Attenuation}\right) \underbrace{\left(\frac{1}{2\gamma}\right)\left(\frac{1}{r+\delta-\mu}\right) x(\kappa_j - \kappa_i)}_{CE_{ij}}.$$

The preceding proposition implies that if all transition rates are equal, there is attenuation bias. Moreover, the attenuation bias is quite severe even if one considers settings with long expected policy durations. For example, suppose the effective discount rate is 5% and that the expected regime life is ten years, with $N = 11$ and $\lambda = .01$. Here attenuation bias is equal to 69% of the causal effect. Consider instead an expected regime life of twenty years, with $N = 11$ and $\lambda = .005$. Although it would be tempting to treat such a long regime as equivalent to permanent, the attenuation bias here amounts to 52% of the causal effect.

4.2. Overshooting and Incorrect Signing of Causal Effects

The settings described in the previous subsection featuring binary assignment or equality of all transition rates are comforting inasmuch as the empiricist can claim treatment responses measured in such environments are conservative estimates of causal effects. Of course, it is not clear why one would want conservative estimates. After all, conservative estimates of elasticities lead to downward bias in the estimated deadweight loss arising from government interventions. Further, magnitudes, as distinct from signs, are often used as the basis for falsifying underlying theories. Finally, the settings analyzed in the previous subsection might be viewed as atypical. After all, it is seldom the case that a policy variable is binary, and it is seldom the case that policy transitions are uniformly distributed.

These arguments notwithstanding, one might hope that the settings considered in the preceding subsection are still instructive regarding the nature of the wedge between treatment responses and causal effects. After all, conventional wisdom holds that agents will simply respond less aggressively to transient government policies, so that attenuation bias may be expected to be a general feature. The following proposition shows that this is not the case.

Proposition 3 *If there are three or more policy states, there exists a continuum of transition rates such that a proper subset of treatment responses are opposite in sign to their respective causal effects. If there are at least three policy states, there exists a continuum of transition rates such* *that a proper subset of treatment responses are equal in sign but larger in absolute value than their respective causal effects.*

We defer for now discussion of the intuition behind the preceding proposition, since it is most readily understood in the context of experiments featuring permanent assignment, the subject of the next subsection.

4.3. Transitions to Absorbing States

Empiricists might be interested in obtaining a heuristic sense of the nature of the bias they face, and the probabilities of each bias form. This subsection provides a practical guide based on settings where the states to which the policy variable may possibly transition ("transition-to states") are absorbing.

To begin, when transition-to states are absorbing, there is a simple relationship between treatment responses and causal effects, with:

$$\begin{aligned} TR_{ij} &= CE_{ij} + \frac{1}{2\gamma}\left[q_i^{**}(x) - q_i(x)\right] \\ q_i^{**}(x) &\equiv \frac{\kappa_i x}{r+\delta-\mu}. \end{aligned} \tag{35}$$

Apparently, if transition-to states are absorbing, the wedge between treatment responses and causal effects is proportional to the wedge between true shadow values accounting for transience (q_i) and the shadow value if the current state were to be permanently assigned (q_i^{**}) as derived in equation (17).

If the current regime is i, with all potential transition-to regimes being absorbing, the Bellman equation is:

$$\begin{aligned} rV^i(s,x) &= \max_a \quad (a-\delta s)V_s^i(s,x) + \mu x V_x^i(s,x) + \frac{1}{2}\sigma^2 x^2 V_{xx}^i(s,x) \\ &\quad + \sum_{j\neq i} \lambda_{ij}[V^j(s,x) - V^i(s,x)] + \kappa_i x s - \psi a - \gamma a^2. \end{aligned} \tag{36}$$

In the preceding equation, the absorbing state value functions (V^j) can be obtained from the constant policy model of Section 2. Solving, we obtain:

$$q_i(x) = q_i^{**}(x) + \frac{x}{(r+\delta-\mu)[1+(r+\delta-\mu)\mathbb{E}(T_i)]}\sum_{j\neq i} P_{ij}(\kappa_j - \kappa_i). \tag{37}$$

From equations (35) and (37) we have the following proposition.

Proposition 4 *If each possible transition-to state is absorbing and the conditional expectation of the policy variable change is zero, then every treatment response is equal to its respective causal effect. If instead the conditional expectation of the policy variable change is not zero, every treatment response is biased, with the ratio of treatment response to causal effect approaching one for policy variable changes large in absolute value and becoming unboundedly large in absolute value for infinitesimal policy variable changes.*

The preceding proposition is useful in that it provides guidance regarding possible safe-harbor conditions an empiricist can invoke to ensure (rough) equality of treatment responses and causal effects. The conjunction of absorbing transition-to states and mean-zero expected changes is one potential safe-harbor assumption. Another safe-harbor assumption is the conjunction of absorbing transition-to states and large policy variable changes. Conversely, the proposition reveals that treatment responses derived from small policy variable changes are especially unreliable as guides to causal effects–absent any notion of rational or irrational response inertia (e.g. Chetty (2012)).

From equations (35) and (37) we have the following proposition delineating bias regions and magnitudes.

Proposition 5 *If each possible transition-to state is absorbing and the conditional expectation of the change in the policy variable is not zero, the sign and magnitude of the treatment response bias depends on the realized assignment category k. If the conditional expectation of the change in κ is positive, then*

$$
\begin{aligned}
\kappa_k - \kappa_i &< 0 \Rightarrow TR_{ik} < CE_{ik} < 0 \\
\kappa_k - \kappa_i &\in \left(0, \frac{\sum_{j \neq i} P_{ij}(\kappa_j - \kappa_i)}{1 + (r + \delta - \mu)\mathbb{E}(T_i)} \right) \Rightarrow TR_{ik} < 0 < CE_{ik} \\
\kappa_k - \kappa_i &\geq \frac{\sum_{j \neq i} P_{ij}(\kappa_j - \kappa_i)}{1 + (r + \delta - \mu)\mathbb{E}(T_i)} \Rightarrow 0 \leq TR_{ik} < CE_{ik}.
\end{aligned}
$$

If the conditional expectation of the change in κ *is negative, then*

$$\begin{aligned}
\kappa_k - \kappa_i &> 0 \Rightarrow TR_{ik} > CE_{ik} > 0 \\
\kappa_k - \kappa_i &\in \left(\frac{\sum_{j \neq i} P_{ij}(\kappa_j - \kappa_i)}{1 + (r + \delta - \mu)\mathbb{E}(T_i)}, 0 \right) \Rightarrow TR_{ik} > 0 > CE_{ik} \\
\kappa_k - \kappa_i &\leq \frac{\sum_{j \neq i} P_{ij}(\kappa_j - \kappa_i)}{1 + (r + \delta - \mu)\mathbb{E}(T_i)} \Rightarrow 0 \geq TR_{ik} > CE_{ik}.
\end{aligned}$$

The intuition for the preceding proposition is as follows. Consider a setting in which the conditional expectation of the change in κ is positive. If the realized treatment is negative ($\kappa_k < \kappa_i$), the treatment response will have the correct negative sign but will overshoot the causal effect in absolute value terms. Objectively bad news becomes very bad news if arriving news was expected to be positive. If instead the realized treatment features a large increase in κ, the treatment response will have the correct positive sign but be biased downwards. This is because if the conditional expectation of the change in κ is positive, firms already invest at a higher rate prior to treatment. Finally, if the realized treatment entails a sufficiently small increase in κ, the treatment response will actually have the wrong sign. Objectively good news becomes bad news if even better news had been expected.

The proposition allows one to readily compute the probability of the various biases, attenuation, overshooting, and sign reversals, depending on the distribution of the treatments, expected regime life, and discount rates. One simply needs to compute the probability of the treatment random variable falling into the bias regions described. In general, the results are not comforting. To illustrate, consider an economy in which $r + \delta - \mu = .05$, with an expected current regime life of five years, and a future permanent change in κ being drawn from a granular set of points uniformly distributed on [0,1]. From the proposition it follows that the treatment response is negative for realized changes less than 0.5/[1+.05(5)]=0.40. That is, evaluated ex ante, the probability of the realized treatment response having the wrong sign is 40%. If the effective discount rate is increased from 5% to 10%, or the expected regime life increases from 5 to 10 years, the sign reversal probability is equal to 33%.

Figure 1 offers a quantitative example illustrating the bias regions for the case of absorbing transition-to states. The figure plots the ratio of treatment response to causal effects for all possible realized policy treatments. The example assumes: the change in κ is uniformly distributed on [-0.20, 0.80]; the expected regime life is five years; and $r + \delta - \mu = .10$.

Consistent with Proposition 5, the figure shows treatment responses overshoot causal effects for realized $\Delta\kappa < 0$. Since the expected change in κ is positive (0.30) bad news becomes very bad news. Treatment responses overshoot causal effects by a very wide margin in the case of small negative treatments. Figure 1 also shows a region where treatment responses are opposite in sign to causal effects. In particular, firms cut investment in reaction to increases in κ on the interval (0, 0.20). The percentage downward bias is extremely large for small increases in κ.

The preceding examples illustrate that sign reversals and overshooting are not anomalies requiring pathological policy generating processes. Rather, one can readily compute bias probabilities depending on the statistical properties of the randomized treatment under consideration. In the preceding example, the probability of the experiment yielding an incorrect sign was 20%, while the probability of overshooting was also 20%. The remaining 60% of potential experimental treatments would result treatment responses attenuated relative to causal effects.

Taken together, the results of this subsection imply that for a natural experiment involving permanent random assignment, the ideal experiment features a mean zero policy variable change. If the policy variable change is not mean zero, bias results. Nevertheless, there is an analytical argument for seeking out experiments in which the realized change in the policy variable is large. Further, if the empiricist is pursuing the more limited goal of estimating a causal effect sign, realized policy changes opposite in sign to expected changes are sign-robust.

4.4. Bias Bounds

This subsection demonstrates some inherent limitations on the nature and severity of biases that can arise in natural policy experiments.

Consider first the problem of attenuation bias. A natural question to ask is whether there exists any policy generating process such that all treatment responses are equal to zero. The answer is no, as shown in the following lemma.

Lemma 1 *There is no set of policy transition rates such that all treatment responses are equal to zero.*

Of course, the preceding lemma provides only a modicum of comfort in relation to the problem of attenuation bias since it still allows for the possibility of severe attenuation. Further, concerning attenuation bias, there is no guarantee that biases cancel if one were to average treatment responses over all possible transitions. Recall, Proposition 1 showed that with binary assignment, both treatment responses are biased downward. And Proposition 2 showed that all possible treatment responses are biased downward by the same proportion if all transition rates are equal ($\lambda_{ij} = \lambda$).

Consider next the problem of treatment responses overshooting causal effects. The following lemma shows that regardless of the assumed transition rates, overshooting cannot possibly occur in the case of worst-to-best state transitions.

Lemma 2 *The treatment response associated with a transition from the worst to best state is always less than its respective causal effect.*

The preceding lemma follows from the fact that with time-varying policies, the shadow value in the best (worst) state cannot be greater than (less than) the shadow value under permanently best (worst) state policy assignments (formula (27)).

Intuition suggests other factors mitigating the extent of treatment response overshooting. To motivate the argument, suppose there are three policy states with low, medium and high values of κ, with the low and high states being absorbing. If the shadow value of capital in the medium state is less than its value under permanent assignment, there will be overshooting in the event of a transition to the high state, but undershooting in the event of a transition to the low state. That is, overshooting for one transition implies undershooting for another transition. The following lemma reveals this to be a more general phenomenon.

Lemma 3 *If there exists a transition from state j to a (better) state $i < j$ with treatment response exceeding its respective causal effect by $k > 0$, then the sum of the treatment responses for transitions from state i to the best state* (1) *and from the worst state* (N) *to j must fall below the sum of their respective causal effects by at least k.*

5. Safe-Harbor Assumptions and a Correction

Section 4 demonstrates random assignment is not generally sufficient for equality of treatment responses and causal effects in dynamic settings. The first objective of this section is to provide empiricists with a set of safe-harbor conditions ensuring that treatment responses are unbiased estimators of causal effects. We next consider conditions for sign-robustness. Finally, we derive corrections for cases in which the described auxiliary assumptions are not satisfied.

Proposition 4 provided one safe-harbor assumption for settings in which all transition-to states are absorbing. In such cases, the necessary and sufficient condition for equality of treatment responses and causal effects is that the conditional expectation of the policy variable change is zero. As shown below, this is actually a special case of a more general proposition.

We begin by deriving a restriction on transition rates such that treatment responses converge to causal effects. To motivate this restriction, it is worth returning to the method used to derive the causal effects implied by the underlying theory. Recall, causal effects (equation (19)) are obtained by differentiating the optimal accumulation function with respect to the government policy variable. In such comparative statics exercises, it is as-if agents initially thought the initial policy was permanent, with God stepping in and perturbing the respective policy variable by an infinitesimal amount, with a credible promise never to perturb the policy variable in the future. Thus, one would expect treatment responses to approximate causal effects if the policy generating process features "almost completely unanticipated" policy changes that are "nearly-permanent." Indeed, we have the following result.

Proposition 6 *As transition rates tend to zero, each treatment response converges to its respective causal effect.*

Given the limited availability of rare-event natural experiments, it would be useful to derive auxiliary assumptions with broader real-world applicability than outlined in the previous proposition. Conveniently, we have the following result.

Proposition 7 *Necessary and sufficient conditions for treatment responses to equal causal effects for each policy transition are that the best and worst policy states are absorbing while remaining states are either absorbing or feature policy variable changes with conditional expectation equal to zero.*

The intuition for the proposition is as follows. Causal effects (comparative statics) measure agent responses to completely unanticipated and permanent changes in policy variables. Therefore, each treatment response is equal to its respective causal effect if the shadow value of capital in every state is equal to its respective shadow value under permanent assignment to that state, despite the policy variable being transient, with transience being the well-spring of experiments. Under the conditions described in Proposition 7, the expected value of the policy variable at each future date is its current value. And if the policy variable is indeed a martingale, the shadow value capitalizes the current policy state as-if lasting into perpetuity. That is, at each instant the firm acts as-if the current policy state will last forever, despite knowing it will not. In such environments the econometrician can directly extract causal effects from treatment responses even if agents are facing highly transient policies.

An empiricist might be concerned that this section's safe-harbor conditions rely upon our initial assumption that adjustment costs take the standard quadratic form γa^2. However, it is apparent that Proposition 6 and Proposition 7 actually hold for the more general class of adjustment cost functions that are convex in accumulation (a). To see this, note that regardless of the particular convex function assumed, the underlying source of bias is the wedge between the true shadow value and the shadow value that would arise under permanent assignment. Under the condition stated in Proposition 6 the shadow value wedge is infinitesimal, and under the conditions stated in Proposition 7 the wedge is zero. We have the following lemma.

Lemma 4 *Suppose adjustment costs are a convex function of accumulation. If all policy transitions are rare events (Proposition 6), then all treatment responses approximate their respective causal effects. If the policy variable is a martingale (Proposition 7), all treatment responses equal their respective causal effects.*

In some cases, an empiricist may only be interested in correctly estimating the sign of a causal effect. A description of safe harbor provisions guaranteeing correct estimation of causal effect signs is therefore worthwhile. We begin by recalling the only type of bias is attenuation bias if there is binary assignment (Proposition 1) or if all transitions rates are equal (Proposition 2). Further, it was shown in Proposition 5 that if the transition-to state is absorbing, the sign of the treatment response is equal in sign to the causal effect for realized treatments opposite in sign to the expected treatment.

We now derive a final safe-harbor provision ensuring sign-robustness. Since optimal accumulation is increasing in the shadow value of capital, sign reversals can be ruled out if q^i is decreasing in i. With this in mind, we use the econometrician's identifying assumption that κ and x are independent stochastic processes in order to rewrite the shadow value (equation (27)) as the discounted marginal product, with the policy variable now indexed by time:

$$q^i(x_0) = \int_0^{\infty} \{\mathbb{E}_0[\kappa_t]\mathbb{E}_0[x_t]\} e^{-(r+\delta)t} dt. \tag{38}$$

Since the current flow return is strictly decreasing in the current state i, it follows from the preceding equation that q^i is strictly decreasing in i if the policy state process is stochastically monotone. Letting the superscript j index the initial state, we recall the continuous-time Markov chain for the policy state ι_t^j is said to be *stochastically monotone* if:

$$\Pr[\iota_t^j \leq k] \geq \Pr[\iota_t^{j+1} \leq k] \quad \forall\ k \in \{1, ..., N\},\ j \in \{1, ..., N-1\},\ t \in \mathbb{R}_+. \tag{39}$$

That is, if the policy state is stochastically monotone, then at all future dates t, ι_t^j is first-order stochastic dominant to ι_t^{j+1}.

Keilson and Kester (1977), derive necessary and sufficient conditions for stochastic monotonicity (Theorem 2.1). Applying their theorem, we have the following proposition.

Proposition 8 *Each treatment response has the same sign as its respective causal effect under policy transition rates such that the following matrix has non-negative off diagonal elements:*

$$\begin{bmatrix} 1 & 0 & ... & 0 \\ 1 & 1 & ... & 0 \\ ... & ... & ... & ... \\ 1 & 1 & 1 & 1 \end{bmatrix}^{-1} \begin{bmatrix} -\sum_{j\neq 1}\lambda_{1j} & \lambda_{12} & ... & \lambda_{1N} \\ \lambda_{21} & -\sum_{j\neq 2}\lambda_{2j} & ... & \lambda_{2N} \\ ... & ... & ... & ... \\ \lambda_{N1} & \lambda_{N2} & ... & -\sum_{j\neq N}\lambda_{Nj} \end{bmatrix} \begin{bmatrix} 1 & 0 & ... & 0 \\ 1 & 1 & ... & 0 \\ ... & ... & ... & ... \\ 1 & 1 & 1 & 1 \end{bmatrix}.$$

Stochastic monotonicity thus represents another possible identifying assumption for the more limited goal of correct estimation of causal effect signs. A number of commonly-utilized processes are stochastically monotone, such as random walks with absorbing or impenetrable boundaries, as shown by Keilson and Kester (1977).

5.2. Mapping Treatment Responses to Causal Effects

This subsection analyzes how to infer causal effects based upon measured treatment responses when the underlying policy generating process does not satisfy the necessary conditions for unbiasedness as stated in Propositions 6 and 7.

The inference problem confronting the econometrician can be viewed as an application of the heterogeneous treatment effects model commonly used in the fields of medicine, education and labor. In particular, since the policymaker does not know the value of firm-specific adjustment cost parameters, she does not know how firms will respond to policy changes. The treatment response of an arbitrary firm m is given by

$$\begin{aligned} TR_{ij}^{m} &= \Theta_m(q_j - q_i) \\ \Theta_m &\equiv \frac{1}{2\gamma_m}. \end{aligned} \tag{40}$$

Heterogeneous effects is a situation that commonly confronts empiricists. Here there is an additional source of complexity since the treatment size is not directly observable. Rather, the size of treatment $(q_j - q_i)$ is determined via the policy transition matrix and equation (29). An unbiased estimator of the mean value of Θ, denoted μ_Θ, is given by the mean treatment response $(\overline{TR}_{ij})$ normalized by size of the treatment. Using equation (29) to compute the relevant state-contingent shadow values, the inferred mean treatment response is:

$$\widehat{\mu}_\Theta = \frac{\overline{TR}_{ij}}{q_j - q_i} = \frac{\overline{TR}_{ij}}{(c_j - c_i)x}. \tag{41}$$

Therefore, the causal effect implied by an observed treatment response is:

$$CE_{ij} = \widehat{\mu}_\Theta(q_j^{**} - q_i^{**}) = \left(\frac{\overline{TR}_{ij}}{q_j - q_i}\right)(q_j^{**} - q_i^{**}) = \left(\frac{\overline{TR}_{ij}}{(c_j - c_i)x}\right)(q_j^{**} - q_i^{**}) \tag{42}$$

with $c_j - c_i$ determined by equation (29).

At this point it is worthwhile to comment on the study by Cummins, Hassett and Hubbard (1994), who use firm responses to imputed changes in the shadow value of capital to infer adjustment cost parameters. Implicit in their shadow value imputation is the assumption that each tax code change is unanticipated and expected to be permanent. They argue their working assumption overstates the change in shadow values and overstates adjustment costs. However, the analysis of Section 4 shows that their imputation method may overstate, understate, or have sign that is opposite to the true change in shadow values. It follows from equation (41) that their inferred adjustment costs can be overstated, understated, or have the wrong sign.

6. The Lucas Critique: Application, Extension and Correction

Up to this point attention has been confined to determining the relationship between measured treatment responses and theory-implied causal effects. This section considers a related but distinct question: whether and how an historic treatment response in a given economy can be used to forecast future treatment responses in that same economy or in different economies.

6.1. External and Internal Validity

A trivial observation that following from the model is that it is not valid, as a general matter, to use the treatment response in one economy to forecast the treatment response in another economy, even if the change in the policy variable is the same and the economies have identical real technologies, e.g. adjustment costs and operating profit functions. To see this, note that it follows from equation (29) that if economies A and B have different policy transition rates, then $\mathbf{T}^A \neq \mathbf{T}^B$, implying differences in shadow values and treatment responses. In other words, as argued by Lucas (1976), if the policy generating process is changed (here represented by $\mathbf{T}$), decision rules (here the TR_{ij}) will change. This is an important limitation on the external validity of treatment responses derived from ideal natural experiments.

In fact, the Lucas Critique is just one limit on extrapolation. The model also illustrates important limits to the *internal validity* of measured treatment responses arising from natural experiments. Specifically, even if one considers a single economy and holds fixed the policy transition matrix $\mathbf{T}$, so that the Lucas Critique does not apply, it is not valid, as a general matter, to use an historic treatment response to forecast future treatment responses.

We illustrate the preceding argument by way of three numerical examples. Each example approximates a different policy environment. All three examples consider variation in the κ policy variable across ten states, with jumps across neighboring states of equal size. Each example assumes that only nearest-neighbor transitions are possible. The effective discount rate is set to 10%.

In the first example it is assumed that the transition rate for an up move is $\lambda_u = .15$ and the transition rate for a down move is $\lambda_d = .05$. The implied expected regime duration is five years. The results are shown in Table 1. As shown, the fact that $\lambda_u > \lambda_d$ results in asymmetric treatment responses. Since an up move is more likely than a down move, shadow values are pushed upward, implying relatively large (small) treatment responses for downward (upward) transitions. For example, starting at $\kappa = 1.9$ the treatment response for an up move is .449, while that for a down move is -.697, a 55% difference. In addition, treatment responses vary across initial states, despite the fact that the transition rates for each state are identical and causal effects are linear. For example, the minimum treatment response is .449 and the maximum treatment response is .978, a difference of 118%.

The second example approximates a polarized political system. As in the preceding example, jumps in κ across neighboring states are of equal size, and only nearest-neighbor transitions are possible. In contrast to the first example, we now assume that for states 1 through 5, "conservatives" are in charge, and there is a bias toward further relaxation of environmental regulations with $\lambda_u = .15$ and $\lambda_d = .05$. Conversely, if the initial state is 6-10, "liberals" are in charge, and there is a bias toward further tightening of regulations with $\lambda_u = .05$ and $\lambda_d = .15$. As in the prior example, the expected policy regime duration is 5 years. As shown in Table 2, here one sees even larger differences between treatment responses across states. For example, the minimum treatment response is .451 and the maximum treatment response is 2.208, a difference of 390%. Further, under this policy generating process the highest treatment responses involve transitions between the interim states 5 and 6.

The final example approximates a political system with a tendency toward centrism. As in the preceding two examples, jumps in κ across neighboring states are of equal size, and only nearest-neighbor transitions are possible. However, it is now assumed that for states 1 through 5 there is a bias toward tightening environmental regulation with $\lambda_d = .15$ and $\lambda_u = .05$. Conversely, for states 6 through 10, there is a bias toward relaxing regulations with $\lambda_d = .05$ and $\lambda_u = .15$. As in the prior two examples, the expected policy regime duration is 5 years. The results are shown in Table 3. Here again one sees large differences in treatment responses across the states. In contrast to the preceding example, the lowest treatment response is for transitions between states 5 and 6, with the highest treatment responses near the boundary.

6.2. Overcoming Limits to External and Internal Validity

A natural question to ask is how the limits to internal and external validity detailed in the previous subsection can be overcome.

Consider first the problem of internal validity. To fix ideas, suppose the empiricist has measured some historic treatment response associated with κ jumping from state i to state j. She is now interested in using this historical evidence to forecast how firms in this same economy will respond to a future transition from from say κ_h to κ_k. From equation (41) we have:

$$\widehat{TR}_{hk} = \widehat{\mu}_{\Theta}(q_k - q_h) = \left(\frac{\overline{TR}_{ij}}{q_j - q_i} \right) (q_k - q_h) = \left[\frac{c_k - c_h}{c_j - c_i} \right] \overline{TR}_{ij}. \tag{43}$$

Extrapolation can be understood as a two-step process. In the first step, one computes the accumulation response per unit of shadow value $\widehat{\mu}_{\Theta}$ by normalizing the historical treatment response by the change in shadow value at that transition date. This is then re-scaled by the change in shadow value associated with the future-date transition. The final term in square-brackets, the ratio of shadow value changes, is computed via equation (29).

Given its apparent brevity, it is worth emphasizing that the last term in squared brackets in equation (43) generally depends upon all transition rates and all possible realizations of the policy variable, not just those in the directly-relevant set $\{h, i, j, k\}$. However, in exceptional cases the preceding formula simplifies. For example, if the auxiliary assumptions described in Propositions 6 or 7 are satisfied, we have:

$$\frac{c_k - c_h}{c_j - c_i} = \frac{\kappa_k - \kappa_h}{\kappa_j - \kappa_i} \Rightarrow \widehat{TR}_{hk} = \left[\frac{\kappa_k - \kappa_h}{\kappa_j - \kappa_i} \right] \overline{TR}_{ij}. \tag{44}$$

That is, if the underlying policy generating process meets the auxiliary assumptions required for equality of treatment responses and causal effects, the extrapolation of treatment responses within a policy generating process requires a simple adjustment for the relative size of changes in the policy variable. Imputations along the lines of formula (44) are commonplace. However, it is apparent that such imputations are not valid as a general matter.

Consider next the problem of external validity (Lucas (1976)). Continuing with the preceding example, suppose the econometrician would like to forecast how firms in Economy B will respond to a change in κ from κ_h^B to κ_k^B based upon the evidence provided by the observed response of firms in Economy A to a change in κ from κ_i^A to κ_j^A. Applying equation (29) we have:

$$\widehat{TR}_{hk}^B = \widehat{\mu}_{\Theta}(q_k^B - q_h^B) = \left(\frac{\overline{TR}_{ij}^A}{q_j^A - q_i^A} \right) (q_k^B - q_h^B) = \left[\frac{c_k^B - c_h^B}{c_j^A - c_i^A} \right] \overline{TR}_{ij}^A. \tag{45}$$

That is, correct extrapolation across two economies requires taking explicit account of their respective policy generating processes via equation (29).

7. Numerical Example: Greenstone (2001)

This section illustrates how one can extract causal effects from observed treatment responses in real-world dynamic settings. We consider the empirical environment analyzed by Greenstone (2001) since this paper is rare in actually providing sufficient detail allowing one to estimate transition rates. The primary emphasis here is not on specific numerical results but rather to outline out the basic method in a real-world setting. Since our model was not tailor-made for the policy setting considered by Greenstone, some approximations are necessary. For example, we set our analysis in continuous-time in order to facilitate analytical solutions whereas the regulations Greenstone considers might be better approximated using a discrete-time setup.

Greenstone (2001) considers the effect of environmental regulations, exploiting quasi-random regulatory assignment arising from the Clean Air Act (CAA). Under the CAA, counties are designated as achieving attainment or non-attainment status in four pollutant categories: Carbon Monoxide, Ozone, Sulfur Dioxide, and Total Suspended Particulates (TSP). Firms emitting the relevant pollutant in non-attainment counties face surcharges and taxes. The Pulp, Paper, Iron and Steel sectors emit all 4 pollutants, implying 16 possible combinations of attainment versus non-attainment status over all pollutants. The Stone, Clay, Glass and Concrete sectors emit Ozone, Sulfur Dioxide and TSP, implying 8 possible combinations of regulatory assignment. The Petroleum Refining sector emits Carbon Monoxide, Ozone and Sulfur Dioxide, implying 8 possible combinations of regulatory assignment. The Non-Ferrous Metals sector emits Carbon Monoxide and Sulfur Dioxide, implying 4 possible combinations of regulatory assignment. Other sectors emit either one pollutant or none.

Table 4 provides detail on the possible regulation categories and assumed κ values. The numerical examples assume each non-attainment infraction results in a reduction in κ. The baseline κ for a fully compliant firm is normalized at 1. The following κ reductions are assumed for non-attainment for each pollutant: Carbon Monoxide=0.30; Ozone=0.25; Sulfur Dioxide=0.20; and TSP=0.10. The lowest possible κ value is 0.15 for a firm emitting all four pollutants and residing in a county that has non-attainment status for all four pollutants.

The key input into the analysis are the transitions rates. Greenstone (2001) reports summary statistics allowing one to infer for each pollutant the conditional probability of starting a given (five-year) observation period in Attainment/Non-Attainment status and ending the period in Attainment/Non-Attainment status. Using the Kolmogorov forward equations, we calibrate the continuous-time transition rates in order to match these conditional probabilities.[3] Table 5 reports transition rates across categories. The average regime life across the 16 possible categories is 1.4 periods or 7 years. We set $r + \delta - \mu = 0.50$ given that each observation period is five years.

Table 6 reports the ratio of treatment response to causal effect for the 12 possible regulation category transitions for firms in the Non-Ferrous Metals sector. As shown, the most common bias is attenuation. However, there are two transitions (16.7%) for which treatment responses overshoot causal effects. Table 7 reports the ratio of treatment response to causal effect for the 56 possible regulation category transitions for firms in the Stone, Clay, Glass and Concrete sectors. The most common bias is attenuation. However, there are two transitions (3.6%) for which treatment response is opposite in sign to the causal effect. And there are eight transitions (14.3%) for which the treatment response overshoots the causal effect. Table 8 reports the ratio of treatment response to causal effect for the 56 possible regulation category transitions for firms in the Petroleum sector. As shown, the most common bias is attenuation. However, there are eight transitions (14.3%) for which treatment response overshoots the causal effect. Finally, Table 9 reports the ratio of treatment response to causal effect for the 240 possible regulation category transitions for firms in the Pulp, Paper, Iron and Steel sectors. Yet again the most common form of bias is attenuation. However, there are 8 transitions (3%) for which treatment response is opposite in sign to the causal effect. And there are 24 transitions (10%) for which treatment response overshoots the causal effect.

Under the data generating process exploited by Greenstone (2001), sign reversals are uncommon. This is because the most common and probable transition is to a neighboring state. In contrast, we recall that sign reversals tend to occur when a large positive or negative change is expected, while the realized change is relatively small. Thus, for those interested in understanding the causal effect sign of environmental regulations, the policy generating process exploited by Greenstone is attractive. However, for those wishing to estimate deadweight losses of future environmental regulations, naïve substitution of his estimates into standard deadweight formulae would generally

[3] See Ross (1996) pages 242-250.

result in a severe understatement of welfare losses.

Setting aside the particular application considered in this section, it is apparent that only by way of opening up the black-box of the policy generating process, as done here, can one hope to understand the sign and magnitude of biases arising in applied empirical exercises in dynamic settings. This can be done by the empiricist or following Greenstone (2001) the empiricist can provide enough detail on the underlying process such that others can perform independent back-of-the-envelope bias estimates.

8. Concluding Remarks

This paper analyzed the relationship between theory-implied causal effects and measured treatment responses in dynamic economies. It was shown that independent policy assignment is insufficient for valid inference. With binary assignment, or uniformly distributed transitions, treatment responses understate causal effects, and substantially so under plausible parameterizations. With more than two policy regimes, treatment responses can understate, overstate, or have a sign that is opposite to causal effects. Moreover, the probability of overshooting and sign reversals can be high under plausible assumptions about the distribution of policy innovations. We also show important limits to the generalizability of historic treatment responses. It is not valid, in general, to extrapolate treatment responses within or across policy generating processes. Together, these results call into question the economic meaning and utilization of a broad range of elasticity estimates shaping policy decisions.

As constructive steps, we derive auxiliary assumptions required to ensure equality of treatment responses and causal effects in dynamic settings. It was shown that if all possible policy changes are rare-events, treatment responses approximate causal effects. However, we offered a more widely-applicable safe-harbor condition, showing that treatment responses equal causal effects if the policy variable is a martingale. In the special case of permanent assignment, each possible treatment response is equal to its respective causal effect if, and only if, the expected policy variable change has mean zero.

In certain cases, empiricists are willing to settle for the more limited objective of correct inference of causal effect signs. We derive a number safe-harbor assumptions that empiricists can invoke as safe-harbors: binary assignment; uniformly distributed policy transitions; stochastic monotonicity of the policy variable; or permanent random assignment opposite in sign to the expected assignment. As final constructive steps, we showed how observed treatment responses can be mapped backed to causal effects, or extrapolated within or across policy generating processes.

It follows from our analysis that an apparent falsification of a given theory based on an incorrect sign prediction may be a false-falsification. In particular, in addition to "*the* identifying assumption" of random assignment, natural policy experiments must be understood as predicated upon equally-important identifying assumptions about policy generating processes. Of course, this is a special example of the arguments of W.V. Quine who pointed out that an empirical falsification only allows one to either reject the underlying theory or to reject an experimenter's auxiliary assumptions. Unfortunately, these auxiliary assumptions have gone unrecognized, unstated and unsatisfied in much of the quasi-experimental literature.

In fact, this paper simply takes the central argument of *Mostly Harmless Econometrics* a necessary step further, illustrating the need to move beyond random assignment. As Angrist and Pischke (2009) state, "The description of an ideal experiment also helps you formulate causal questions precisely. The mechanics of an ideal experiment highlight the forces you'd like to manipulate and the factors you'd like to hold constant." In the view of many, random assignment constitutes this ideal. Instead, we have shown that random assignment is properly seen as just one ingredient in the making of an ideal policy experiment in dynamic settings.

Appendix: Proofs and Derivations

For brevity, let:

$$\begin{aligned} R &\equiv r + \delta - \mu \\ c &\equiv \frac{\kappa}{r + \delta - \mu}. \end{aligned}$$

Growth Option Value with Constant Government Policies

Since the terms in the Bellman equation scaled by s cancel each other, satisfaction of the equation demands the growth option value satisfies the following ODE:

$$rG(x) = \mu x G_x(x) + \frac{1}{2}\sigma^2 x^2 G_{xx}(x) + \underbrace{(\Gamma c^2)x^2 - 2\Gamma c\psi x + \Gamma\psi^2}_{Dividend}. \tag{46}$$

It follows from the preceding equation and the Feynman-Kac formula that the growth option value G is equal to the value of a claim to a dividend stream that is linear-quadratic in x. To value this linear-quadratic claim, conjecture a linear-quadratic value function G with unknown constants and substitute this function into the preceding differential equation. One obtains:

$$G(x) = \frac{1}{4\gamma}\left[\left(\frac{c^2}{r - 2\mu - \sigma^2}\right)x^2 - \left(\frac{2c\psi}{r-\mu}\right)x + \frac{\psi^2}{r}\right]. \tag{47}$$

Growth Option Value with Regime Shifts

To derive the growth option value, return to the Bellman system (25) and confine attention to the remaining terms that have not been zeroed out, those terms not scaled by s. We have the following system pinning down growth option value:

$$\begin{aligned} \left(r + \sum_{j\neq 1}\lambda_{1j}\right)G^1(x) &= \mu x G^1_x(x) + \frac{1}{2}\sigma^2 x^2 G^1_{xx}(x) + \sum_{j\neq 1}\lambda_{1j}G^j(x) + \Gamma c_1^2 x^2 - 2\Gamma c_1\psi x + \Gamma\psi^2 \quad (48) \\ &\;\;... \\ \left(r + \sum_{j\neq N}\lambda_{Nj}\right)G^N(x) &= \mu x G^N_x(x) + \frac{1}{2}\sigma^2 x^2 G^N_{xx}(x) + \sum_{j\neq N}\lambda_{Nj}G^j(x) + [\Gamma c_N^2]x^2 - 2\Gamma c_N\psi x + \Gamma\psi^2. \end{aligned}$$

We conjecture a growth option value that is linear-quadratic in x, with regime shifts. The following lemma follows directly by substituting the conjectured linear-quadratic solution into the system of ODEs in equation (48).

Lemma: *The no-bubbles solution to the differential equations*

$$\left(r+\sum_{j\neq 1}\lambda_{1j}\right)J^1(x) \;=\; \mu x J^1_x(x)+\frac{1}{2}\sigma^2x^2J^1_{xx}(x)+\sum_{j\neq 1}\lambda_{1j}J^j(x)+\phi_1x^2+\widetilde{\phi}_1x+\widehat{\phi} \tag{49}$$
$$\ldots$$
$$\left(r+\sum_{j\neq N}\lambda_{Nj}\right)J^N(x) \;=\; \mu x J^N_x(x)+\frac{1}{2}\sigma^2x^2J^N_{xx}(x)+\sum_{j\neq N}\lambda_{Nj}J^j(x)+\phi_Nx^2+\widetilde{\phi}_Nx+\widehat{\phi}$$

is

$$\begin{bmatrix} J^1(x) \\ \ldots \\ J^N(x) \end{bmatrix} = x^2[\mathbf{T}(r-2\mu-\sigma^2)]^{-1}\begin{bmatrix} \phi_1 \\ \ldots \\ \phi_N \end{bmatrix} + x[\mathbf{T}(r-\mu)]^{-1}\begin{bmatrix} \widetilde{\phi}_1 \\ \ldots \\ \widetilde{\phi}_N \end{bmatrix} + \frac{\widehat{\phi}}{r}\begin{bmatrix} 1 \\ \ldots \\ 1 \end{bmatrix}. \tag{50}$$

Proposition 2

We have the following $N \times N$ augmented transition matrix:

$$\mathbf{T}(R) \equiv \begin{bmatrix} R+(N-1)\lambda & -\lambda & \ldots & -\lambda \\ -\lambda & R+(N-1)\lambda & \ldots & -\lambda \\ \ldots & \ldots & \ldots & \ldots \\ -\lambda & -\lambda & \ldots & R+(N-1)\lambda \end{bmatrix} \Rightarrow \mathbf{T}^{-1} = \frac{1}{R(R+N\lambda)}\begin{bmatrix} R+\lambda & \lambda & \ldots & \lambda \\ \lambda & R+\lambda & \ldots & \lambda \\ \ldots & \ldots & \ldots & \ldots \\ \lambda & \lambda & \ldots & R+\lambda \end{bmatrix} \tag{51}$$

The shadow values are:

$$\begin{bmatrix} q^1(x) \\ q^2(x) \\ \ldots \\ q^N(x) \end{bmatrix} = \frac{x}{R(R+N\lambda)}\begin{bmatrix} R+\lambda & \lambda & \ldots & \lambda \\ \lambda & R+\lambda & & \lambda \\ \ldots & \ldots & \ldots & \ldots \\ \lambda & \lambda & \ldots & R+\lambda \end{bmatrix}\begin{bmatrix} \kappa_1 \\ \kappa_2 \\ \ldots \\ \kappa_N \end{bmatrix}. \tag{52}$$

Taking differences across rows, the treatment response of the shadow value is:

$$q^j(x)-q^i(x) = \left[1-\frac{N\lambda}{R+N\lambda}\right]\frac{1}{R}(\kappa_j-\kappa_i)x. \tag{53}$$

And the result follows from $\Delta a = \Delta q/2\gamma$.■

Proposition 3

Suppose there are $N \geq 3$ states and assume there are three states (l, m, h) that are absorbing as a system in that once the policy process enters any one of these three states, the state never transitions to a state $i \notin \{l, m, h\}$. Transition rates amongst states outside $\{l, m, h\}$ and into $\{l, m, h\}$ can be set arbitrarily. In this case, the shadow value in any of these three states can be computed by focusing on the system confined to $\{l, m, h\}$. To fix notation, assume the states vary in the intensity of government intervention, with:

$$\kappa_h < \kappa_m < \kappa_l.$$

To prove the first part of the lemma it is sufficient to find transition rates for the three state system with sign reversals. To this end, assume the medium and low states are absorbing. We have:

$$\mathbf{T} = \begin{bmatrix} R & 0 & 0 \\ 0 & R & 0 \\ -\lambda_{hl} & -\lambda_{hm} & R + \lambda_{hl} + \lambda_{hm} \end{bmatrix} \Rightarrow \mathbf{T}^{-1} = \frac{1}{R} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ \frac{\lambda_{hl}}{R+\lambda_{hl}+\lambda_{hm}} & \frac{\lambda_{hm}}{R+\lambda_{hl}+\lambda_{hm}} & \frac{R}{R+\lambda_{hl}+\lambda_{hm}} \end{bmatrix}. \quad (54)$$

This implies:

$$\begin{aligned} q^m(x) &= \left(\frac{x}{R}\right) \kappa_m \\ q^h(x) &= \frac{x}{R\,(R + \lambda_{hl} + \lambda_{hm})} \left[\lambda_{hl}\kappa_l + \lambda_{hm}\kappa_m + R\kappa_h\right]. \end{aligned} \quad (55)$$

It follows that

$$\lambda_{hl} > \frac{R(\kappa_m - \kappa_h)}{\kappa_l - \kappa_m} \Rightarrow q^h(x) > q^m(x).$$

To prove the second part of the lemma it is sufficient to find transition rates for the three state system with overshooting. To this end, assume that now the low and high states are absorbing. The transition matrix is:

$$\mathbf{T} = \begin{bmatrix} R & 0 & 0 \\ -\lambda_{ml} & R + \lambda_{ml} + \lambda_{mh} & -\lambda_{mh} \\ 0 & 0 & R \end{bmatrix} \Rightarrow \mathbf{T}^{-1} = \frac{1}{R} \begin{bmatrix} 1 & 0 & 0 \\ \frac{\lambda_{ml}}{R+\lambda_{ml}+\lambda_{mh}} & \frac{R}{R+\lambda_{ml}+\lambda_{mh}} & \frac{\lambda_{mh}}{R+\lambda_{ml}+\lambda_{mh}} \\ 0 & 0 & 1 \end{bmatrix}. \quad (56)$$

We have:

$$\begin{bmatrix} q^l(x) \\ q^m(x) \\ q^h(x) \end{bmatrix} = \frac{x}{R} \begin{bmatrix} 1 & 0 & 0 \\ \frac{\lambda_{ml}}{R+\lambda_{ml}+\lambda_{mh}} & \frac{R}{R+\lambda_{ml}+\lambda_{mh}} & \frac{\lambda_{mh}}{R+\lambda_{ml}+\lambda_{mh}} \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \kappa_l \\ \kappa_m \\ \kappa_h \end{bmatrix}. \quad (57)$$

We are interested in conditions such that the treatment response exceeds the causal effect for a transition from the medium to low state:

$$q^l(x) - q^m(x) > \frac{x(\kappa_l - \kappa_m)}{R}. \tag{58}$$

The preceding condition is met if:

$$\lambda_{ml}(\kappa_l - \kappa_m) < \lambda_{mh}(\kappa_m - \kappa_h).\blacksquare$$

Lemma 1

Consider policy experiments involving changes in κ. The treatment response is zero across all transitions only if the shadow value is constant across regimes. This holds only if there is some constant k such that the vector c=$k\mathbf{1}$. But this implies the existence of an augmented transition matrix $\mathbf{T}$ satisfying:

$$k\mathbf{1} = \mathbf{T}^{-1}(R)\begin{bmatrix}\kappa_1\\ ...\\ \kappa_N\end{bmatrix} \Rightarrow \mathbf{T}(R)\mathbf{1} = \begin{bmatrix}R\\ ...\\ R\end{bmatrix} = \frac{1}{k}\begin{bmatrix}\kappa_1\\ ...\\ \kappa_N\end{bmatrix}. \tag{59}$$

This is a contradiction.$\blacksquare$

Lemma 3

Let stars denote shadow values under permanent policies. From Lemma 2 if follows:

$$q^1(x) - q^N(x) \leq q^1_*(x) - q^N_*(x). \tag{60}$$

This inequality may be rewritten as:

$$[q^1(x)-q^i(x)]+[q^j(x)-q^N(x)]+[q^i(x)-q^j(x)] \leq [q^1_*(x)-q^i_*(x)]+[q^j_*(x)-q^N_*(x)]+[q^i_*(x)-q^j_*(x)]. \tag{61}$$

Rearranging terms, it follows:

$$[q^1(x)-q^i(x)]+[q^j(x)-q^N(x)] \leq [q^1_*(x)-q^i_*(x)]+[q^j_*(x)-q^N_*(x)]-\{[q^i(x)-q^j(x)]-[q^i_*(x)-q^j_*(x)]\}. \tag{62}$$

Under the conditions stipulated in the lemma, it follows:

$$[q^1(x) - q^i(x)] + [q^j(x) - q^N(x)] \leq [q^1_*(x) - q^i_*(x)] + [q^j_*(x) - q^N_*(x)] - 2\gamma k. \tag{63}$$

The result then follows by dividing the preceding equation by 2γ.■

Proposition 6

As each transition rate approaches zero, $\mathbf{T}^{-1}$ approaches $R^{-1}\mathbf{1}$ implying each q^i approaches $R^{-1}\kappa_i$ implying each TR_{ij} approaches CE_{ij}.■

Proposition 7

Consider an array of distinct policies 1 to N with the indexing convention being that the shadow value under permanent policies is decreasing in the index. A necessary and sufficient condition for each treatment response to equal its respective causal effect is that the difference between state-contingent shadow values is equal to shadow values under permanent policies. Thus, there must exist some k such that for all i:

$$q^i(x) = \frac{\kappa_i x}{r+\delta-\mu} + k. \tag{64}$$

Therefore, we must identify conditions such that the following equilibrium conditions can be met under this functional form:

$$\begin{aligned}
\left(r+\delta+\sum_{j\neq 1}\lambda_{1j}\right) q^1(x) &= \mu x q^1_x(x) + \frac{1}{2}\sigma^2 x^2 q^1_{xx}(x) + \sum_{j\neq 1}\lambda_{1j} q^j(x) + \kappa_1 x \qquad (65)\\
&\cdots\\
\left(r+\delta+\sum_{j\neq N}\lambda_{Nj}\right) q^N(x) &= \mu x q^N_x(x) + \frac{1}{2}\sigma^2 x^2 q^N_{xx}(x) + \sum_{j\neq N}\lambda_{Nj} q^j(x) + \kappa_N x.
\end{aligned}$$

Substituting in the required functional form and canceling terms we obtain:

$$\begin{aligned}
(r+\delta)\,k &= \sum_{j\neq 1}\lambda_{1j}\left[\frac{(\kappa_j-\kappa_1)x}{r+\delta-\mu}\right] \qquad (66)\\
(r+\delta)\,k &= \sum_{j\neq 2}\lambda_{2j}\left[\frac{(\kappa_j-\kappa_2)x}{r+\delta-\mu}\right]\\
&\cdots\\
(r+\delta)\,k &= \sum_{j\neq N-1}\lambda_{N-1j}\left[\frac{(\kappa_j-\kappa_{N-1})x}{r+\delta-\mu}\right]\\
(r+\delta)\,k &= \sum_{j\neq N}\lambda_{Nj}\left[\frac{(\kappa_j-\kappa_N)x}{r+\delta-\mu}\right].
\end{aligned}$$

First note that any solution to the preceding system entails $k=0$ since the right-side of the first equation is non-positive while the right-side of the last equation is non-negative. It follows that any candidate solution to the system entails:

$$\begin{aligned}\lambda_{1j} &= 0 : j = 2, ..., N \\ \lambda_{Nj} &= 0 : j = 1, ..., N-1.\end{aligned}$$

Further, it must be the case that:

$$0 = \sum_{j \neq i} \lambda_{ij} \left[\frac{(\kappa_j - \kappa_i)x}{r + \delta - \mu} \right] : i = 2, ..., N-1 \tag{67}$$

The preceding equation can be rewritten as:

$$\sum_{j \neq i} \left(\frac{\lambda_{ij}}{\sum_{j \neq i} \lambda_{ij}} \right) (\kappa_j - \kappa_i) = 0 : i = 2, ..., N-1. \blacksquare \tag{68}$$

References

[1] Aaron, Henry J., 1990, Lessons for Tax Reform, in J. Slemrod Ed., *Do Taxes Matter? The Impact of the Tax Reform Act of 1986,* MIT Press.

[2] Abel, Andrew, 1983, Tax Neutrality in the Presence of Adjustment Costs, *Quarterly Journal of Economics*, 705-712.

[3] Abel, Andrew and Janice Eberly, 1997, An Exact Solution for the Investment and Value of a Firm Facing Uncertainty, Adjustment Costs, and Irreversibility, *Journal of Economic Dynamics and Control* (21), 831-852.

[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] Atanasov, Vladimir and Bernard Black, 2015, Shock-Based Causal Inference in Corporate Finance Research, ECGI Working Paper 448/2015.

[7] Auerbach, Alan J., 1986, Tax Reform and Adjustment Costs: The Impact on Investment and Market Value, *International Economic Review* (30), 831-852.

[8] Auerbach, Alan J. and James R. Hines, Jr., 1988, Investment Tax Incentives and Frequent Tax Reforms, *American Economic Review* (78), 211-216.

[9] Banerjee, Abhijit, Esther Duflo, Rachel Glennerster, and Cynthia Kinnan, 2015, The Miracle of Microfinance? Evidence from a Randomized Evaluation, *American Economic Journal: Applied Economics*, 22-53.

[10] Chay, Kenneth Y. and Michael Greenstone, 2005, Does Air Quality Matter? Evidence from the Housing Market, Journal of Political Economy (113), 376-424.

[11] Chetty, Raj, 2012, Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply, *Econometrica* (80), 969-1018.

[12] Crepon, Bruce, Florencia Devoto, Esther Duflo, and William Pariente, 2013, Estimating the Impact of Microcredit on Those Who Take it Up: Evidence from a Randomized Experiment, 2015, *American Economic Journal: Applied Economics*, 123-150.

[13] Cummins, Jason G., Kevin A. Hassett, and R. Glenn Hubbard, 1994, A Reconsideration of Investment Behavior Using Tax Reforms as Natural Experiments, *Brookings Papers on Economic Activity* (2), 1-74.

[14] Deschenes, Olivier and Michael Greenstone, 2007, Economic Impacts of Climate Change: Evidence from Random Weather Fluctuations, *American Economic Review.*

[15] Gan, Jie, 2007, Collateral, Debt Capacity and Corporate Investment: Evidence from a Natural Experiment, *Journal of Financial Economics* (85), 709-754.

[16] Gourio, Francois and Jianjun Miao, 2008, Transitional Dynamics of Dividend and Capital Gains Tax Cuts, *Review of Economic Dynamics* (14), 368-383.

[17] Greenstone, Michael, 2002, The Impacts of Environmental Regulations on Industrial Activity: Evidence from the 1970 and 1977 Clean Air Act Amendments and Census of Manufacturers, *Journal of Political Economy* (110).

[18] Hassett, Kevin A. and Gilbert Metcalf, 1999, Investment with Uncertain Tax Policy: Does Random Taxation Discourage Investment, *Economic Journal* (109), 372-393.

[19] Keilson, Julian and Adri Kester, 1977, Monotone Matrices and Monotone Markov Processes, *Stochastic Processes and Their Applications* (5), 231-241.

[20] Kolokoltsov, Vassili N., 2011 Stochastic Monotonicity and Duality for One-Dimensional Markov Processes, *Mathematical Notes* (89), 652-666.

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

[22] Romer, Christina and David Romer, 2010, Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks, *American Economic Review*, 763-801.

[23] Slemrod, Joel, 1990, The Economic Impact of the Tax Reform Act of 1986, in J. Slemrod Ed., *Do Taxes Matter? The Impact of the Tax Reform Act of 1986*, MIT Press: Cambridge.

[24] Slemrod, Joel, 1992, Do Taxes Matter? Lessons from the 1980's, *American Economic Review* (82), 250-256.

Table 1: Treatment Responses Constant Transition Rates

Initial Kappa	Transition Up	Transition Down
2.0	NA	-0.449
1.9	0.449	-0.697
1.8	0.697	-0.833
1.7	0.833	-0.908
1.6	0.908	-0.950
1.5	0.950	-0.971
1.4	0.971	-0.978
1.3	0.978	-0.958
1.2	0.958	-0.813
1.1	0.813	NA

Table 2: Treatment Responses Polarized Political System

Initial Kappa	Transition Up	Transition Down
2.0	NA	-0.451
1.9	0.451	-0.705
1.8	0.705	-0.875
1.7	0.875	-1.140
1.6	1.140	-2.208
1.5	2.208	-1.140
1.4	1.140	-0.875
1.3	0.875	-0.705
1.2	0.705	-0.451
1.1	0.451	NA

Table 3: Treatment Responses Centrist Political System

Initial Kappa	Transition Up	Transition Down
2.0	NA	-0.768
1.9	0.768	-0.871
1.8	0.871	-0.818
1.7	0.818	-0.680
1.6	0.680	-0.420
1.5	0.420	-0.680
1.4	0.680	-0.818
1.3	0.818	-0.870
1.2	0.870	-0.769
1.1	0.769	NA

Table 4: Regulatory Attainment Categories and Profitability Factors

Numerical Category	Carbon Monoxide	Ozone	Sulfur Dioxide	Total Suspended Particulates	Profitabiilty Factor
1	Attain	Attain	Attain	Attain	1.00
2	Attain	Attain	Attain	Non	0.85
3	Attain	Attain	Non	Attain	0.80
4	Attain	Non	Attain	Attain	0.75
5	Non	Attain	Attain	Attain	0.70
6	Attain	Attain	Non	Non	0.65
7	Attain	Non	Attain	Non	0.60
8	Attain	Non	Non	Attain	0.55
9	Non	Attain	Attain	Non	0.55
10	Non	Attain	Non	Attain	0.50
11	Non	Non	Attain	Attain	0.45
12	Attain	Non	Non	Non	0.40
13	Non	Attain	Non	Non	0.35
14	Non	Non	Attain	Non	0.30
15	Non	Non	Non	Attain	0.25
16	Non	Non	Non	Non	0.10

Table 5: Transition Rate Matrix

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1	0.0000	0.0671	0.0182	0.0946	0.0238	0.0000	0.0000	0.0001	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
2	0.6463	0.0000	0.0000	0.0001	0.0000	0.0181	0.0944	0.0000	0.0238	0.0000	0.0001	0.0002	0.0000	0.0000	0.0000	0.0000
3	0.7403	0.0000	0.0000	0.0000	0.0000	0.0671	0.0000	0.0944	0.0001	0.0238	0.0000	0.0002	0.0000	0.0000	0.0001	0.0000
4	0.1725	0.0000	0.0000	0.0000	0.0000	0.0000	0.0670	0.0182	0.0000	0.0000	0.0239	0.0000	0.0000	0.0000	0.0000	0.0000
5	0.2821	0.0000	0.0000	0.0000	0.0000	0.0000	0.0001	0.0000	0.0671	0.0181	0.0944	0.0000	0.0000	0.0001	0.0001	0.0000
6	0.0000	0.7405	0.6464	0.0000	0.0000	0.0000	0.0001	0.0002	0.0000	0.0000	0.0000	0.0944	0.0241	0.0002	0.0001	0.0000
7	0.0000	0.1727	0.0001	0.6464	0.0000	0.0000	0.0000	0.0000	0.0002	0.0000	0.0000	0.0182	0.0000	0.0240	0.0000	0.0000
8	0.0001	0.0001	0.1725	0.7402	0.0000	0.0000	0.0001	0.0000	0.0001	0.0000	0.0000	0.0672	0.0001	0.0000	0.0239	0.0000
9	0.0000	0.2822	0.0000	0.0000	0.6462	0.0000	0.0000	0.0000	0.0000	0.0001	0.0000	0.0000	0.0180	0.0944	0.0001	0.0000
10	0.0000	0.0000	0.2822	0.0002	0.7405	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0002	0.0671	0.0002	0.0946	0.0000
11	0.0001	0.0000	0.0000	0.2821	0.1725	0.0000	0.0001	0.0000	0.0000	0.0000	0.0000	0.0000	0.0001	0.0670	0.0182	0.0000
12	0.0001	0.0000	0.0000	0.0000	0.0000	0.1729	0.7409	0.6468	0.0001	0.0001	0.0000	0.0000	0.0000	0.0001	0.0001	0.0239
13	0.0002	0.0000	0.0000	0.0000	0.0000	0.2824	0.0000	0.0002	0.7401	0.6462	0.0003	0.0000	0.0000	0.0000	0.0000	0.0945
14	0.0001	0.0000	0.0000	0.0000	0.0000	0.0002	0.2821	0.0000	0.1726	0.0001	0.6463	0.0000	0.0000	0.0000	0.0000	0.0182
15	0.0000	0.0000	0.0002	0.0000	0.0000	0.0000	0.0000	0.2822	0.0003	0.1727	0.7403	0.0002	0.0000	0.0001	0.0000	0.0671
16	0.0000	0.0004	0.0003	0.0004	0.0003	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.2827	0.1730	0.7409	0.6468	0.0000

Table 6: Ratio of Treatment Response to Causal Effect Nonferrous Metals

CATEGORY	1	3	5	10
1	NA	0.3972	0.6203	0.5310
3	0.3972	NA	1.0663	0.6202
5	0.6203	1.0663	NA	0.3972
10	0.2655	0.2326	0.1135	NA

Table 7: Ratio of Treatment Response to Causal Effect Stone, Clay, Glass, and Concrete

CATEGORY	1	2	3	4	6	7	8	12
1	NA	0.4123	0.3965	0.6517	0.4030	0.5620	0.5386	0.5070
2	0.4123	NA	0.3492	1.0109	0.3961	0.6518	0.6017	0.5385
3	0.3965	0.3492	NA	1.6725	0.4117	0.7274	0.6522	0.5622
4	0.6517	1.0109	1.6725	NA	-0.2188	0.4123	0.3972	0.4036
6	0.4030	0.3961	0.4117	-0.2188	NA	1.6745	1.0131	0.6525
7	0.5620	0.6518	0.7274	0.4123	1.6745	NA	0.3517	0.3970
8	0.5386	0.6017	0.6522	0.3972	1.0131	0.3517	NA	0.4121
12	0.5070	0.5385	0.5622	0.4036	0.6525	0.3970	0.4121	NA

Table 8: Ratio of Treatment Response to Causal Effect Petroleum Refining

CATEGORY	1	3	4	5	8	10	11	15
1	NA	0.3980	0.6523	0.6205	0.5390	0.5315	0.6346	0.5716
3	0.3980	NA	1.6693	1.0655	0.6518	0.6205	0.7697	0.6347
4	0.6523	1.6693	NA	0.4617	0.3975	0.4107	0.6198	0.5312
5	0.6205	1.0655	0.4617	NA	0.3761	0.3980	0.6514	0.5390
8	0.5390	0.6518	0.3975	0.3761	NA	0.4638	1.0644	0.6204
10	0.5315	0.6205	0.4107	0.3980	0.4638	NA	1.6650	0.6517
11	0.6346	0.7697	0.6198	0.6514	1.0644	1.6650	NA	0.3984
15	0.5716	0.6347	0.5312	0.5390	0.6204	0.6517	0.3984	NA

Table 9: Ratio of Treatment Response to Causal Effect
Pulp, Paper, Iron and Steel

CATEGORY	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
1	NA	0.4131	0.3966	0.6522	0.6202	0.4038	0.5622	0.5388	0.5510	0.5308	0.6348	0.5069	0.5031	0.5870	0.5712	0.5440
2	0.4131	NA	0.3471	1.0109	0.8273	0.3968	0.6517	0.6017	0.6199	0.5813	0.7179	0.5382	0.5301	0.6345	0.6107	0.5702
3	0.3966	0.3471	NA	1.6747	1.0674	0.4134	0.7279	0.6526	0.6745	0.6203	0.7709	0.5621	0.5504	0.6632	0.6347	0.5861
4	0.6522	1.0109	1.6747	NA	0.4601	-0.2172	0.4122	0.3970	0.4245	0.4094	0.6203	0.4031	0.4098	0.5508	0.5307	0.5024
5	0.6202	0.8273	1.0674	0.4601	NA	-0.8946	0.3883	0.3760	0.4126	0.3968	0.6523	0.3936	0.4027	0.5622	0.5385	0.5059
6	0.4038	0.3968	0.4134	-0.2172	-0.8946	NA	1.6712	1.0113	1.0662	0.8272	1.0390	0.6512	0.6189	0.7703	0.7176	0.6332
7	0.5622	0.6517	0.7279	0.4122	0.3883	1.6712	NA	0.3514	0.4611	0.4053	0.8283	0.3962	0.4084	0.6201	0.5814	0.5294
8	0.5388	0.6017	0.6526	0.3970	0.3760	1.0113	0.3514	NA	NA	0.4591	1.0667	0.4112	0.4226	0.6739	0.6198	0.5492
9	0.5510	0.6199	0.6745	0.4245	0.4126	1.0662	0.4611	NA	NA	0.3494	1.0119	0.3746	0.3952	0.6519	0.6015	0.5370
10	0.5308	0.5813	0.6203	0.4094	0.3968	0.8272	0.4053	0.4591	0.3494	NA	1.6744	0.3872	0.4105	0.7276	0.6519	0.5604
11	0.6348	0.7179	0.7709	0.6203	0.6523	1.0390	0.8283	1.0667	1.0119	1.6744	NA	-0.8999	-0.2214	0.4120	0.3963	0.4013
12	0.5069	0.5382	0.5621	0.4031	0.3936	0.6512	0.3962	0.4112	0.3746	0.3872	-0.8999	NA	0.4571	1.0679	0.8283	0.6182
13	0.5031	0.5301	0.5504	0.4098	0.4027	0.6189	0.4084	0.4226	0.3952	0.4105	-0.2214	0.4571	NA	1.6788	1.0140	0.6504
14	0.5870	0.6345	0.6632	0.5508	0.5622	0.7703	0.6201	0.6739	0.6519	0.7276	0.4120	1.0679	1.6788	NA	0.3492	0.3933
15	0.5712	0.6107	0.6347	0.5307	0.5385	0.7176	0.5814	0.6198	0.6015	0.6519	0.3963	0.8283	1.0140	0.3492	NA	0.4080
16	0.5440	0.5702	0.5861	0.5024	0.5059	0.6332	0.5294	0.5492	0.5370	0.6500	0.4013	0.6182	0.6504	0.3933	0.4080	NA

Figure 1: Treatment Response to Causal Effect Ratio