Learning From Natural Policy Experiments When Agents Do Too*

Christopher A. Hennessy†

November 2017

Abstract

This paper considers the correct interpretation and extrapolation of evidence coming from sequences of seemingly-ideal exogenous policy shocks when the underlying data generating process is not known to either agents or econometricians. Firms accumulate a stock variable facing exogenous shocks driven by a N-state continuous-time Markov chain, with unobserved arrival intensity. As shown, learning gives rise to "causal parameter drift" even with constant a data generating process. The correct intrepretation of shock responses hinges upon the exact time pattern of realized shocks, as well as (generally unstated) parametric assumptions about priors and potential arrival intensities. Conveniently, closed-form formulae are given for mapping observed shock responses back to theory-implied causal effects (comparative statics).

*Preliminary and Incomplete: Please do not quote or cite without author permission.

†London Business School (CEPR and ECGI). Contact: chennessy@london.edu.

1 Introduction

In their influential textbook, *Mostly Harmless Econometrics*, Angrist and Pischke (2009) argue empirical evidence derived from exogenous shocks represents a stand-alone product: "A principle that guides our discussion is that most of the estimators in common use have a simple interpretation that is not heavily model dependent." The asserted model-light, and assumption-light, nature of their proposed tool-kit may go a long way to explaining its increasing popularity relative to structural models (e.g. Hennessy and Whited (2005, 2007) which feature explicit assumptions about underlying technologies and stochastic processes.

Angrist and Pischke go on to argue that "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." In Angrist and Pischke (2010), they go on to herald the search for sources of exogenous variation as a "credibility revolution" in the causal effect identification problem, going on to chide macroeconomists and macroeonometricians, in particular, for failing to share in their revolutionary zeal. In his response, Sims (2010) argues that microeconometricians have just as much, or more, to learn from macroeconometricians, but fails to specify exactly what that means at an operational level.

Through a conjunction of repetition and routinization, the original inference problem, causal effect identification, has been converted into the problem of demonstrating that exploited shocks are orthogonal. To wit, the question "What is your identification strategy?" is generally treated as being equivalent to "Why should we believe the exploited shocks are as-good-as-random?" Effectively, the means, orthogonal shocks, have replaced the ends, empirical estimation of the sign and magnitude of causal effects. As pointed out by Deaton (2010), part of the reason that means have been converted into ends is that empiricists frequently fail to provide up-front a formal definition of the exact object they are attempting to credibly identify. To paraphrase Deaton in his critique of instrumental variables, "This goes well beyond the problem of looking under the streetlamp. Here we have a problem where we have some control over where the light falls but instead proclaim that wherever the light falls is the thing we wanted to see."

The ostensible goal of much empirical work is to empirically estimate signs and magnitudes of causal effects. As argued by Heckman (2010), causal effects are well-defined within the context of formal theoretical models, since they correspond to comparative statics. Indeed, as pointed out by Athey and Milgrom (1999), when theorists flush out the empirical predictions of their proposed theories, their workhorse methodology is comparative statics analysis. The basic question addressed in this paper is: What is the relationship between the objects measured by empiricists, responses to orthogonal policy shocks, and those of interest to theorists, as well as those using elasticities to perform welfare analysis, comparative statics.

Effectively, in performing comparative statics, one compares outcomes in one parameterized economy to outcomes in an otherwise identical economy endowed with a different parameterization. Phrased differently, in performing comparative statics in the policy sphere, one acts as-if economic agents had posited a perpetually fixed value for the policy parameter, with the government then changing the parameter in a way that was completely unanticipated, with agents then expecting the government's policy change will also be its last policy change. Importantly, this type of formal analysis seems to accord with the informal notion of causality invoked by Angrist and Pischke (2009) at the start of their textbook: "A causal relationship is useful for making predictions about the consequences of changing circumstances or policies; it tells us what would happen in alternative (or 'counterfactual') worlds."

In reality, empiricists seldom exploits shocks that are completely unanticipated. After all, if shocks were completely unanticipated, they would be one in a million year events. In reality, legislated tax and regulatory changes are much more frequent, to say nothing of passive quasi-random changes in the real effective burdens imposed by taxation and regulation due to changes in the rate of inflation, as well as changes in both judicial interpretation and executive branch implementation. Second, and perhaps more importantly, if a measured policy shock were completely permanent, there would be little reason to study how agents respond to the shock. After all, the primary motivations for analyzing policy shock responses is to forecast future shock responses or for the government to re-optimize that policy variable in the future. It is apparent that the primary motivations for studying shock responses are logically inconsistent with the oft-invoked assumption that a measured shock was expected to be permanent.

Hennessy and Strebulaev (2017) demonstrate that with realistic real-world policy transience, there is generally a large wedge between measured shock responses and theory-implied causal effects, with sign reversals also being a real possibility. However, they also go on to show that if the underlying stochastic process generating shocks is known, then shock responses can readily be translated to infer what they imply about theory-implied causal effects. Heuristically, one must undo the effect of the true policy transition matrix featuring transience, and then impose the effects of a different policy transition matrix, one featuring unanticipated and permanent shocks (infinitesimal shock arrival rates). An analogous technique is proposed for extrapolating shock responses across data generating processes, so that cross-economy empirical estimates can be compared.

One frequently heard response to such proposed quasi-structural approaches to inference is: "We do not even know the true underlying data generating process–and firms probably do not know it either." It is this response that constitutes the motivation for the present paper. In particular, the formal contribution of this paper is to extend the framework of Hennessy and Strebulaev (2017) to account for uncertainty and learning about the underlying data generating process. Anticipating, the results of the analysis are both destructive and constructive. First, the destructive contribution of this paper is to show that the challenges to causal effect inference highlighted by Hennessy and Strebulaev (2017) may be just the tip of the iceberg, with learning giving rise to estimator drift. Specifically, with learning, measured shock responses vary over time along with beliefs, and this is true even if the underlying data generating process is fixed. This greatly complicates the causal effect inference, forcing the analyst to make even more parametric assumptions, e.g. assumptions regarding the set of possible data generating processes posited by agents (and econometricians). Second, the constructive contribution of the paper is to illustrate across a variety of dynamic learning environments how shock responses can be extrapolated or mapped back to theory-implied causal effects (comparative statics).

To fix ideas, I consider an econometrician seeking to empirically estimate the causal effect implied by a canonical dynamic theory, the q-theory of capital accumulation under quadratic adjustment costs. Despite focusing on capital, the model describes the behavior of shadow values more generally, with shadow values predicting optimal accumulation of a broad range of stock variables, e.g. debt, cash, inventory, employees, or reputation. Unlike the econometrician, I know the true value of the causal effect parameter. In contrast, the econometrician instead relies upon measured shock responses for inference. Importantly, endogeneity is not an issue, since, by construction, the policy process is orthogonal. In particular, the policy variable is governed by an independent N-state continuous-time Markov chain. Departing from Hennessy and Strebulaev (2017), the true instantaneous arrival rate of shocks is not observed by the firms, nor by the econometricians who study them. That is, I follow Sargent in imposing rational expectations' "communism of models." In this setting, I examine the relationship between *shock responses* (empirically measured responses to policy variable changes) and *causal effects* (Marshallian comparative statics/differentials).

I turn now to related literature. A varied empirical literature exploits (quasi) random assignment in search of causal effects in dynamic economies. Greenstone (2002) and Greenstone and Chay (2005) exploit random air quality regulation, evaluating effects on real 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 Hubbard (1994), treats tax code changes as natural experiments. Romer and Romer (2010) identify exogenous tax changes to estimate their macroeconomic effects. Card and Krueger (1994) analyze minimum wage effects using difference-in-differences. 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.

In spirit, my paper draws upon Lucas (1976) who pointed out the perils of extrapolating macroeconometric evidence across policy rules. In contrast, my paper highlights the inadequacy of the random assignment assumption taken in isolation. Naturally, Lucas does not analyze the relationship between shock responses under dynamic random assignment and causal effects, nor derive formulae for correction and extrapolation.

My paper shares with Keane and Wolpin (2002) the notion that one must account for dynamics and uncertainty in order to correctly infer causal effects. However, there are numerous important differences. First, in terms of context, they analyze a granular dynamic model of contraceptive use and welfare participation. I instead offer a more general/abstract analysis of the effect of dynamics and uncertainty on shadow values, the key determinant of optimal accumulation of a broad array of commonly-studied stock variables. Second, they offer numerical solutions featuring polynomial approximations while I present fully analytic solutions amenable to direct analysis and back-of-the-envelope adjustments. Finally, and most importantly, I consider the problem of causal effect inference in an economy in which agents and econometricians are endowed with symmetric ignorance regarding the data generating process.

Abel (1982) analyzes effects of permanent versus temporary tax policies with no policy variable uncertainty. Specifically, he assumes, in potential violation of rational expectations, that the policy change is completely unanticipated, with policy being binary and deterministic after the complete surprise policy change. Incorporating rational expectations, as well as the probability distribution of policy variables, is essential for our results. House and Shapiro (2008) also ignore uncertainty and present response approximations assuming arbitrarily short-lived policy changes. In their setting, the shadow value of capital is left unchanged by policy changes. In contrast, I analyze the behavior of shadow values, and optimal accumulation, allowing for policy variable uncertainty and expected regime lives of arbitrary duration–in a setting with learning.

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.

I use a general structural model to understand and correct empirical estimates derived from random assignment in dynamic settings. Recovery of a deep structural adjustment cost parameter is an interim step. Cummins, Hassett and Hubbard (1994) estimate adjustment cost parameters based on tax experiments. They abstract from learning and assume each tax change is a complete surprise and viewed as permanent.

My paper is related to but distinct from the literature contesting the external validity of estimates derived from instrumental variables, e.g. Heckman (1997) and Deaton (2010). To begin, in the setting I consider, there is absolutely no instrumentation and by construction no possibility of self-selection, so the oft-discussed issues surrounding the correct interpretation of LATE are extraneous. Second, issues concerning heterogeneous treatment effects are also extraneous to my analysis which assumes agents are homogeneous, or equivalently, considers inference on a firm-by-firm basis.

2 Constant Data Generating Process

This section considers the interpretation of shock responses in the context of an economy with a constant, yet latent, data generating process.

2.1 Technology

Time is continuous and the horizon is infinite. All agents in the economy are risk-neutral and discount cash flows at rate $r > 0$. The firm accumulates a stock variable, treated here as capital, which decays at a constant rate δ. The capital stock evolves as:

$$dK_t = (I_t - \delta K_t)dt. \tag{1}$$

The cash flow of the firm is equal to operating profit, which is linear in K, less the costs of its investment program inclusive of quadratic adjustment costs. Thus, at all points in time, cash flow is described by

$$\Pi_t K_t - [I_t + \gamma I_t^2]. \tag{2}$$

The econometrician is interested in estimating the causal effect of a change in government policy. In the model, government policy operates by changing the profitability factor Π. For example, one can think of government policy affecting profitability through taxation and regulation. To fix ideas, let us assume the econometrician knows all the deep causal parameters aside from the stock adjustment cost parameter γ. From the perspective of estimating the causal effect of government policy, this is the keep parameter in the economy since it determines the responsiveness of firms, as described below.

To fix ideas, let us assume that in contrast to the econometrician, we know the deep structural parameter γ. This will allow us to assess the ability of the econometrician to infer the theory-implied causal effect based upon shock responses. The theory-implied causal effect (CE) is simply the comparative static effect of a change in Π. Effectively, comparative statics tell us the response of firms to completely unanticipated policy changes that are permanent. With permanent government policies, the shadow value of capital is

$$q^{**} = \frac{\Pi}{r + \delta}. \tag{3}$$

The optimal investment policy must equate the shadow value of a unit of capital with the marginal cost of investment. We then have:

$$q = 1 + 2\gamma I^* \Rightarrow I^* = \frac{q - 1}{2\gamma}. \tag{4}$$

From the preceding two equations we obtain the following theory-implied causal effect, which is linear:

$$CE \equiv \frac{\partial I^*}{\partial \Pi} = \frac{(r + \delta)^{-1}}{2\gamma}. \tag{5}$$

From the preceding equation it is apparent that in the present parable economy, inferring causal effects is isomorphic to inferring the adjustment cost parameter γ.

Let us assume that the econometrician attempts to infer causal effects from shock responses. The present section considers a simple shock generating process. In particular, the profitability factor is assumed to be driven by a continuous-time Markov chain with a finite number of states S. Conditional upon a shock taking place, the probability of transitioning from state s to state $\widetilde{s}$ is given by $\rho_{s\widetilde{s}}$ with each such conditional transition probability assumed to be constant over time. Let us assume further, that at all points in time, the instantaneous probability of a policy shock is $\widetilde{\lambda} \in \{0, \lambda\}$. This implies that the expected life of the current state is $1/\widetilde{\lambda}$. Critically, let us assume that the true value of $\widetilde{\lambda}$ is not known by any agent in this economy. That is, neither the econometrician nor the firms know the true value of $\widetilde{\lambda}$.

2.2 Post-Shock Behavior

In the technological setting considered in this section, learning takes a rather simple form. In particular, when and if the first shock takes place, all agents learn that they occupy an economy in which $\widetilde{\lambda} = \lambda$. That is, after the first shock occurs (if it ever does) all agents, firms and the econometrician know the true data generating process. The remainder of this section characterizes firm behavior in this setting (which is analyzed exhaustively in Hennessy and Strebulaev (2017)).

The value of the firm in policy state s is denoted V_s. Accounting for regime changes, and concomitant capital gains, we have the following system of S Hamilton-Jacobi-Bellman equations:

$$\begin{aligned} rV_1(K) &= \max_I \quad (I - \delta K)V_1'(K) + \lambda \sum_{s \neq 1} \rho_{1s}[V_s(K) - V_1(K)] + \Pi_1 K - I - \gamma I^2 \\ &\quad ... \\ rV_S(K) &= \max_I \quad (I - \delta K)V_S'(K) + \lambda \sum_{s \neq S} \rho_{Ss}[V_s(K) - V_S(K)] + \Pi_S K - I - \gamma I^2. \end{aligned} \tag{6}$$

From the preceding HJB equations it follows that the optimal state-contingent investment policy solves:

$$I_s \in \arg\max_I \quad IV_s'(K) - I - \gamma I^2; \; s = 1, ..., S. \tag{7}$$

The solution to the Bellman system (6) has the following functional form:

$$V_1(K) = KQ_1 + G_1; ...; V_S(K) = KQ_S + G_S. \tag{8}$$

It follows that the optimal state-contingent investment policy is:

$$I_s^* = \frac{Q_s - 1}{2\gamma}; \ \ s = 1, ..., S. \tag{9}$$

From equation (9) it follows investment will exhibit a jump each time there is a policy shock. These shock responses are observed by the econometrician and are the basis for her causal inference.

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

$$I_s^* Q_s - I_s^* - \gamma I_s^{*2} = \frac{(Q_s - 1)^2}{4\gamma}. \tag{10}$$

Substituting the accumulation gain from equation (10) and the conjectured value functions into the original system of Bellman equations, we can rewrite the Bellman system as:

$$(r + \delta + \lambda) K Q_1 + (r + \lambda) G_1 = \lambda \sum_{s \neq 1} \rho_{1s} [K Q_s + G_s] + \Pi_1 K + \frac{(Q_1 - 1)^2}{4\gamma} \tag{11}$$

$$...$$

$$(r + \delta + \lambda) K Q_S + (r + \lambda) G_S = \lambda \sum_{s \neq S} \rho_{Ss} [K Q_s + G_s] + \Pi_S K + \frac{(Q_S - 1)^2}{4\gamma}$$

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

$$\begin{aligned} (r + \delta + \lambda) Q_1 &= \lambda \sum_{s \neq 1} \rho_{1s} Q_s + \Pi_1 \\ &... \\ (r + \delta + \lambda) Q_S &= \lambda \sum_{s \neq S} \rho_{Ss} Q_s + \Pi_S \end{aligned} \tag{12}$$

Solving the system of linear equations (12), it follows that the shadow value of capital after the first observed policy shock is:

$$\begin{bmatrix} Q_1 \\ ... \\ Q_S \end{bmatrix} = \mathbf{T}^{-1} \begin{bmatrix} \Pi_1 \\ ... \\ \Pi_S \end{bmatrix} \tag{13}$$

where $\mathbf{T}$

$$\mathbf{T} \equiv \begin{bmatrix} r+\delta+\lambda & -\lambda\rho_{12} & ... & -\lambda\rho_{1N} \\ -\lambda\rho_{21} & r+\delta+\lambda & ... & -\lambda\rho_{2N} \\ ... & ... & ... & ... \\ -\lambda\rho_{N1} & -\lambda\rho_{N2} & ... & r+\delta+\lambda \end{bmatrix}. \tag{14}$$

The derivation of the growth option value function is omitted in the interest of brevity. Essentially, this derivation would involve solving the linear system implied by the terms in the Bellman equation that are not scaled by K.

Consider finally the shock responses that will be observed by the econometrician. The shock response associated with a transition from state s to $\widetilde{s}$ is denoted $SR_{s\widetilde{s}}$. We have:

$$\begin{aligned} SR_{s\widetilde{s}} &\equiv I^*_{\widetilde{s}} - I^*_s \\ &= \frac{1}{2\gamma}(Q_{\widetilde{s}} - Q_s). \end{aligned} \tag{15}$$

We this have the following proposition (Hennessy and Strebulaev (2017)) present the discrete-time analog.

Proposition 1 *With the stipulated constant data generating process, after the first policy shock, all shock responses are time-invariant and can be mapped to their corresponding theory-implied causal effect via:*

$$CE_{s\widetilde{s}} = SR_{s\widetilde{s}} \times \frac{(\Pi_{\widetilde{s}} - \Pi_s)/(r+\delta)}{Q_{\widetilde{s}} - Q_s}$$

where $\mathbf{Q} = \mathbf{T}^{-1}\mathbf{\Pi}$.

As a general matter, shock responses will not be equal to theory-implied causal effects (comparative statics). As shown by Hennessy and Strebulaev (2017), with a known data generating process, there are two special cases in which shock responses will be equal to causal effects: permanent unanticipated shocks or if the profit factor is a martingale.

2.3 Pre-Shock Behavior

Prior to the arrival of any policy shock, agents must engage in Bayesian updating. With this in mind, let

$$B_t \equiv \Pr\left[\widetilde{\lambda} = 0 | \mathcal{F}_t\right]. \tag{16}$$

Conditional upon no policy shock, Bayes' law implies the following evolution of beliefs:

$$B + dB = \frac{B}{B + (1 - \lambda dt)(1 - B)}. \tag{17}$$

Rearranging terms, we obtain the following law of motion for beliefs prior to the arrival of any policy shock:

$$dB = \frac{\lambda B(1 - B)dt}{B + (1 - \lambda dt)(1 - B)}. \tag{18}$$

To rule out trivial cases, for the remainder of the analysis we assume that agents enter the economy with non-dogmatic priors, with $B_0 \in (0, 1)$. After all, if date 0 beliefs are at either extreme, 0 or 1, then it is apparent from the preceding equation that agents never revise their beliefs.

The HJB equation here can be derived heuristically as follows. Let v denote firm value prior to the first policy shock, and assume that, for comparability with the prior subsections results, that the economy starts in state s. That is, prior to the first shock $\Pi = \Pi_s$. We then have the following equilibrium condition stating that expected holding return is just equal to what would be earned if funds were instead invested in the risk-free asset:

$$\begin{aligned}(1 + rdt)v(K, B) \quad = \quad & \max_I \quad [v(K, B) + v_B(K, B)dB + v_K(K, B)dK][B + (1 - B)(1 - \lambda dt)] \\ & + \sum_{\widetilde{s} \neq s} \rho_{s\widetilde{s}} V_s(K)(1 - B)\lambda dt + \left[\Pi_s K - (I + \gamma I^2)\right] dt.\end{aligned} \tag{19}$$

Using the laws of motion of capital and beliefs, and dropping terms of order less than we obtain the following Bellman equation:

$$\begin{aligned}[r + \lambda(1 - B)]v(K, B) \quad = \quad & \max_I \quad \lambda B(1 - B)v_B(K, B) + (I - \delta K)v_K(K, B) \\ & + \lambda(1 - B) \sum_{\widetilde{s} \neq s} \rho_{s\widetilde{s}} [KQ_s + G_s] + \Pi_s K - I - \gamma I^2\end{aligned} \tag{20}$$

The value function prior to the first policy shock takes the following separable form:

$$v(K, B) = Kq(B) + g(B). \tag{21}$$

Substituting the preceding value function into the Bellman equation we obtain:

$$\begin{aligned} [r + \lambda(1 - B)][q(B)K + G(B)] &= \max_I \quad \lambda B(1 - B)[Kq'(B) + g'(B)] + (I - \delta K)q(B) \\ &\quad +\lambda(1 - B)\sum_{\widetilde{s}\neq s} \rho_{s\widetilde{s}}\,[KQ_s + G_s] + \Pi_s K - I - \gamma I^2 \end{aligned} \tag{22}$$

Isolating those terms in the Bellman equation involving the investment policy, it is apparent that the optimal policy is identical in form to that after the first shock. In particular,

$$\begin{aligned} I_s^* &= \frac{q_s - 1}{2\gamma};\ s = 1, ..., S \\ &\Rightarrow\ I_s^* q_s - I_s^* - \gamma I_s^{*2} = \frac{(q_s - 1)^2}{4\gamma}. \end{aligned} \tag{23}$$

Substituting the preceding accumulation gain into the Bellman equation we obtain:

$$\begin{aligned} [r + \lambda(1 - B)][q(B)K + G(B)] &= \lambda B(1 - B)[Kq'(B) + g'(B)] - \delta Kq(B) \\ &\quad +\lambda(1 - B)\sum_{\widetilde{s}\neq s} \rho_{s\widetilde{s}}\,[KQ_{\widetilde{s}} + G_s] + \Pi_s K + \frac{(q_s - 1)^2}{4\gamma}. \end{aligned} \tag{24}$$

The terms in the previous equation scaled by the capital stock must equate. Exploiting this fact, we obtain a linear first-order ODE for the value of a unit of installed capital:

$$[r + \delta + \lambda(1 - B)]q(B) = \lambda B(1 - B)q'(B) + \lambda(1 - B)\sum_{\widetilde{s}\neq s} \rho_{s\widetilde{s}} Q_{\widetilde{s}} + \Pi_s. \tag{25}$$

A derivation of the solution to this ODE is presented in the appendix. We then have the following proposition.

Proposition 2 *With the stipulated constant data generating process, before the first policy shock, the belief that there is no shock possible evolves according to*

$$dB = \frac{\lambda B(1 - B)dt}{B + (1 - \lambda dt)(1 - B)}.$$

The state$-s$ shadow value of capital before the first policy shock is

$$q_s(B) = \frac{\Pi_s}{r+\delta} + \left(\frac{(1-B)\lambda}{r+\delta+\lambda}\right)\left(\sum_{\widetilde{s}\neq s}\rho_{s\widetilde{s}}Q_{\widetilde{s}} - \frac{\Pi_s}{r+\delta}\right).$$

The response to the first policy shock can be mapped to its corresponding theory-implied causal effect via:

$$CE_{s\widetilde{s}} = SR_{s\widetilde{s}} \times \frac{(\Pi_{\widetilde{s}} - \Pi_s)/(r+\delta)}{Q_{\widetilde{s}} - q_s(B)}$$

where $\mathbf{Q} = \mathbf{T}^{-1}\mathbf{\Pi}$.

The expression for the shadow value of capital given in the proposition is intuitive. The first term in large brackets represents the price of a claim paying one dollar in the event of a policy shock arriving. The second term measures the per unit capital gain accruing to the holder of the capital.

A number of insights are apparent from the preceding proposition. First, it is readily verified that, despite the policy generating process being constant, the response to the first policy shock, say a transition from Π_s to $\Pi_{\widetilde{s}}$ will differ from the response to future policy shocks featuring a transition from Π_s to $\Pi_{\widetilde{s}}$. That is, we have causal parameter drift in the sense that learning causes shock responses to vary over time even though the true data generating process is constant over time. In particular, we have the following equations describing shock responses for the first shock versus future shocks.

$$\begin{aligned}
\text{First} \quad &: \quad SR_{s\widetilde{s}} = \frac{1}{2\gamma}\left[Q_{\widetilde{s}} - q_s(B)\right]. \\
\text{Future} \quad &: \quad \frac{1}{2\gamma}(Q_{\widetilde{s}} - Q_s).
\end{aligned}$$

Second, it is apparent that in order to make a correct comparison between causal parameter estimates across time, and in order to correctly map the initial shock response back to the theory-implied causal effect, agent beliefs must be correctly stipulated. Third, and similarly, it is apparent that the magnitude of the first shock response will be a function of the waiting time for the natural experiment. In particular, the longer the waiting time for a natural experiment, the higher is the belief that no shock will arrive. For example, as the waiting time approaches infinity, B will converge to 1.

Finally, and related to this argument, it is apparent that the notion of a policy shock being "unanticipated" is not sufficiently precise to allow for correct inference, extrapolation, and comparisons across shock responses. For example, one possible way in which a first experiment one can be said to be unexpected is that λ is extremely low while the belief that shocks are at all possible $(1-B)$ is high. Another way in which the first experiment can be said to be unexpected is that the belief that shocks are at all possible $(1-B)$ is extremely low while λ is high. But notice, even if the initial instantaneous shock probabilities $(1-B)\lambda$ are identical across these two settings, future shock responses will differ fundamentally since the future instantaneous shock probability, post-learning, is given by λ. In the former scenario, even future policy shocks will come as surprises. In the latter scenario, future policy shocks are expected after the occurrence of the first shock.

3 Heterogeneous Experiments

The previous section considered a constant data generating process. However, in many applied settings it will be the case that the data generating process will change over time. The objective of this section is to flush out the issues and complexities in inference that arise in such settings.

This section considers the following experimental setting. There are potentially $N \geq 2$ sequential policy changes over time. Below, the superscript n will be used to index firm value functions during the time interval leading up to the n-th policy shock. The arrival intensity of shock n is either λ_l^n and λ_h^n, with $\lambda_h^n > \lambda_l^n \geq 0$. In addition to the arrival intensities differing across the experimental rounds, the transition probabilities conditional upon a shock taking place are also allowed to differ across the experimental rounds. In particular, conditional upon shock number n taking place, the probability of transitioning from the current state s to state $\widetilde{s}$ is given by $\rho_{s\widetilde{s}}^n$. Finally, at the start of experimental round $n+1$ (just after shock n but before shock $n+1$) agents start with an exogenous prior belief β^{n+1} that the arrival intensity is low. For example, one might think that prior beliefs will differ across experimental rounds due to the nature of the political process or learning-by-doing on the part of the government.

Consider the beliefs of agents in anticipation of the arrival of shock n. Conditional upon no shock taking place, Bayes' law implies the following evolution of beliefs:

$$B + dB = \frac{(1-\lambda_l^n dt)B}{(1-\lambda_l^n dt)B + (1-\lambda_h^n dt)(1-B)}. \tag{26}$$

Solving the preceding equation we obtain the following law of motion for beliefs:

$$dB = \frac{(\lambda_h^n - \lambda_l^n)B(1-B)dt}{(1-\lambda_l^n dt)B + (1-\lambda_h^n dt)(1-B)} \tag{27}$$

We have then have following equilibrium condition, again with the superscript indexing the upcoming experiment number and with subscripts indexing the current economic state:

$$\begin{aligned}(1+rdt)V_s^n &= \max_I [V_s^n(K,B) + \frac{\partial}{\partial B}V_s^n(K,B)dB + \frac{\partial}{\partial K}V_s^n(K,B)dK][(1-\lambda_l^n dt)B + (1-\lambda_h^n dt)(1-B)] \\ &\quad + [B\lambda_l^n dt + (1-B)\lambda_h^n dt]\sum_{\tilde{s}\neq s}\rho_{s\tilde{s}}V_{\tilde{s}}^{n+1}[K,\beta^{n+1}] + \left[\Pi_s K - (I+\gamma I^2)\right]dt \end{aligned} \tag{28}$$

Expanding the terms dB and dK in the preceding equation and dropping those tending to zero faster than dt, we obtain the following Bellman equation:

$$\begin{aligned}[r + \lambda_l^n B + \lambda_h^n(1-B)]V_s^n &= \max_I (\lambda_h^n - \lambda_l^n)B(1-B)\frac{\partial}{\partial B}V_s^n(K,B) + \frac{\partial}{\partial K}V_s^n(K,B)(I-\delta K) \\ &\quad + [B\lambda_l^n + (1-B)\lambda_h^n]\sum_{\tilde{s}\neq s}\rho_{s\tilde{s}}V_{\tilde{s}}^{n+1}[K,\beta^{n+1}] + \left[\Pi_s K - (I+\gamma I^2)\right] \end{aligned} \tag{29}$$

All value functions will take the separable form:

$$V_s^n(K,B) = Kq_s^n(B) + G_s^n(B). \tag{30}$$

Substituting the value function expressions into the Bellman equation we obtain:

$$\begin{aligned}[r + \lambda_l^n B + \lambda_h^n(1-B)][Kq_s^n(B) + G_s^n(B)] &= \max_I (\lambda_h^n - \lambda_l^n)B(1-B)\frac{\partial}{\partial B}[Kq_s^n(B) + G_s^n(B)] + q_s^n(B)(I-\delta K) \\ &\quad + [B\lambda_l^n + (1-B)\lambda_h^n]\sum_{\tilde{s}\neq s}\rho_{s\tilde{s}}[Kq_{\tilde{s}}^{n+1}(\beta^{n+1}) + G_{\tilde{s}}^{n+1}(\beta^{n+1})] \\ &\quad + \left[\Pi_s K - (I+\gamma I^2)\right]. \end{aligned} \tag{31}$$

Isolating those terms in the Bellman equation involving the investment policy, it is apparent that the optimal policy is:

$$
\begin{aligned}
I_s^{n*} &= \frac{q_s^n(B)-1}{2\gamma};\ s=1,...,S;\ n=1,...,N \\
&\Rightarrow\ I_s^{n*}q_s^n - I_s^{n*} - \gamma\,(I_s^{n*})^2 = \frac{[q_s^n(B)-1]^2}{4\gamma}.
\end{aligned} \tag{32}
$$

Substituting the preceding accumulation gain into the Bellman equation we obtain:

$$
\begin{aligned}
& [r+\delta+\lambda_l^n B+\lambda_h^n(1-B)]Kq_s^n(B)+[r+\lambda_l^n B+\lambda_h^n(1-B)]G_s^n(B)] \\
= \ & \Pi_s K+\frac{[q_s^n(B)-1]^2}{4\gamma}+(\lambda_h^n-\lambda_l^n)B(1-B)\frac{\partial}{\partial B}[Kq_s^n(B)+G_s^n(B)] \\
& +[B\lambda_l^n+(1-B)\lambda_h^n]\sum_{\widetilde{s}\neq s}\rho_{s\widetilde{s}}[Kq_{\widetilde{s}}^{n+1}(\beta^{n+1})+G_{\widetilde{s}}^{n+1}(\beta^{n+1})].
\end{aligned} \tag{33}
$$

Next, note that the terms in the Bellman equation scaled by K must equate. From this fact, we obtain the following ODE for the shadow value of capital:

$$
\begin{aligned}
& [r+\delta+\lambda_l^n B+\lambda_h^n(1-B)]q_s^n(B) \\
= \ & \Pi_s+(\lambda_h^n-\lambda_l^n)B(1-B)\frac{\partial}{\partial B}q_s^n(B)+[B\lambda_l^n+(1-B)\lambda_h^n]\sum_{\widetilde{s}\neq s}\rho_{s\widetilde{s}}q_{\widetilde{s}}^{n+1}(\beta^{n+1}).
\end{aligned} \tag{34}
$$

The preceding linear first-order ODE is solved in the appendix. We then have the following proposition.

Proposition 3 *Suppose there are a finite number of heterogeneous experiments. Prior to the arrival of shock of shock n beliefs evolve according to*

$$
dB=\frac{(\lambda_h^n-\lambda_l^n)B(1-B)dt}{(1-\lambda_l^n dt)B+(1-\lambda_h^n dt)(1-B)}.
$$

The shadow value of capital prior to shock n is given by

$$
q_s^n(B)=\frac{\Pi_s}{r+\delta}+\left[\sum_{\widetilde{s}}\rho_{s\widetilde{s}}^1 q_{\widetilde{s}}^{n+1}(\beta^{n+1})-\frac{\Pi_s}{r+\delta}\right]\left[\frac{[B\lambda_l^n+(1-B)\lambda_h^n](r+\delta)+\lambda_l^n\lambda_h^n}{(r+\delta+\lambda_l^n)(r+\delta+\lambda_h^n)}\right] \tag{35}
$$

The response to policy shock n can be mapped to its corresponding theory-implied causal effect via:

$$
CE_{s\widetilde{s}}=SR_{s\widetilde{s}}\times\frac{(\Pi_{\widetilde{s}}-\Pi_s)/(r+\delta)}{q_{\widetilde{s}}^{n+1}(\beta^{n+1})-q_s^n(B)}
$$

Intuitively, the preceding proposition informs us that the shadow value of capital in any given experimental round is equal to the shadow value of capital under fixed policies plus the expected capital gain associated with a shock (first term in large squared brackets) times the price of a primitive claim paying one dollar in the event of a policy shock arriving. Operationally, the proposition gives one a simple formula that can be applied using backward induction to derive the shadow value of capital at all experimental states. With these shadow values in-hand, the final statement in the proposition allows one to extract the theory-implied causal effect from any observed shock reponse.

Of course, along with its explicit constructive message, there is an implicit destructive message provided by the proposition. After all, the proposition lays bare the fact that making sense of a given observed shock response, extracting its economic meaning, requires making numerous assumptions about beliefs over the set of possible models–itself a high dimensional object. Moreover, the proposition lays bare the fact that shock responses are contingent upon-waiting times. Or, phrased differently, competing empirical estimates are only directly comparable if the underlying data generating process *and* the realization of shock waiting times are identical.

4 Heterogeneous Experiments with Memory

The prior section considered a general finite sequence of heterogeneous experiments. Specifically, the experiments differed in the sense that the shock arrival intensities and the transition-to probabilities were allowed to differ across the experimental rounds. Despite the generality of the setting, a simple analytical characterization was possible. A key feature of that setting was that the shock arrival intensities were assumed to be independent across rounds. While such an assumption may be appropriate in some experimental settings, there may be other settings in which arrival intensities across rounds are correlated. The objective of this section is to examine causal parameter inference in such settings. In particular, we consider the problem of inference across two heterogeneous experiments in which a relatively low (high) shock arrival intensity in the first experimental round implies a relatively low (high) shock arrival intensity in the second experimental rounds. Anticipating, this setting becomes considerably more complex because the beliefs held by agents in the first round carry over to the second experimental round.

It is convenient to solve via backward induction. To begin, we conjecture the value functions take the separable form:

$$\begin{aligned} V_s^1(K,B) &= q_s^1(B)K + G_s^1(B) \\ V_{\widetilde{s}}^2(K,\widetilde{B}) &= q_{\widetilde{s}}^2(\widetilde{B})K + G_{\widetilde{s}}^2(\widetilde{B}) \end{aligned} \tag{36}$$

Consider the second experimental round. In that round, the final round, any shock that occurs is permanent. Thus, shadow value of a unit of installed capital conditional upon the second shock arriving is given by

$$\widehat{q}_{\widetilde{s}}^2 = \frac{1}{r+\delta}\sum_{\widehat{s}} \rho_{\widetilde{s}\widehat{s}}^2 \Pi_{\widehat{s}}.$$

Substituting the preceding expression for the continuation shadow value into equation (**??**), it follows that during the second experimental round, the time between the first and second exogenous shocks, the shadow value of capital is:

$$\begin{aligned} q_{\widetilde{s}}^2(\widetilde{B}) &= \frac{\Pi_{\widetilde{s}}}{r+\delta} + H(\widetilde{B})\sum_{\widehat{s}} \rho_{\widetilde{s}\widehat{s}}^2 (\Pi_{\widehat{s}} - \Pi_{\widetilde{s}}) \\ H(\widetilde{B}) &\equiv \frac{\widetilde{B}\lambda_l(r+\delta) + (1-\widetilde{B})\lambda_h(r+\delta) + \lambda_h\lambda_l}{(r+\delta)(r+\delta+\lambda_l)(r+\delta+\lambda_h)}. \end{aligned} \tag{37}$$

Consider next the shadow value of capital prior to the arrival of the first exogenous shock. Unfortunately, the previous section's analysis cannot be applied here because the continuation shadow value depends upon the revised beliefs that will be formed with the arrival of the first shock. In particular, upon the arrival of the first shock, agents will form the following belief regarding the probability of their being a low arrival intensity of shocks (which they will carry with them as beliefs into the second experimental round):

$$\widetilde{B}(B) = \frac{\lambda_l B}{B\lambda_l + (1-B)\lambda_h}$$

The continuation value conditional upon the first exogenous shock just having arrived is given by:

$$\widehat{V}_s^1(K,B) \equiv \sum_{\widetilde{s}} \rho^1_{s\widetilde{s}} V^2_{\widetilde{s}}[K, \widetilde{B}(B)].$$

We then have the following Bellman equation during the time prior to the first policy shock:

$$\begin{aligned}
[r+\lambda_l B+\lambda_h(1-B)]V_s^1 &= (\lambda_h-\lambda_l)B(1-B)\frac{\partial}{\partial B}V_s^1(K,B)+(I-\delta K)\frac{\partial}{\partial K}V_s^1(K,B) \\
&\quad +[B\lambda_l+(1-B)\lambda_h]\widehat{V}_s^1(K,B)+\left[\Pi_s K-(I+\gamma I^2)\right]. \qquad (38)
\end{aligned}$$

Substituting the conjectured solutions into the Bellman equation we obtain:

$$\begin{aligned}
& [r+\lambda_l B+\lambda_h(1-B)][q_s^1(B)K+G_s^1(B)] \\
= \; & (\lambda_h-\lambda_l)B(1-B)\frac{\partial}{\partial B}[q_s^1(B)K+G_s^1(B)]+(I-\delta K)q_s^1(B) \\
& +[B\lambda_l+(1-B)\lambda_h]\left[K\sum_{\widetilde{s}}\rho^1_{s\widetilde{s}}q^2_{\widetilde{s}}[\widetilde{B}(B)]+\sum_{\widetilde{s}}\rho^1_{s\widetilde{s}}G^2_{\widetilde{s}}[\widetilde{B}(B)]\right]+\left[\Pi_s K-(I+\gamma I^2)\right]
\end{aligned}$$

The terms in the previous equation scaled by the capital stock must equate. Exploiting this fact we obtain an ODE for the value of a unit of installed capital:

$$[r+\delta+\lambda_l B+\lambda_h(1-B)]q_s^1(B)-(\lambda_h-\lambda_l)B(1-B)\frac{\partial}{\partial B}q_s^1(B)-\Pi_s = [B\lambda_l+(1-B)\lambda_h]\sum_{\widetilde{s}}\rho^1_{s\widetilde{s}}q^2_{\widetilde{s}}[\widetilde{B}(B)]. \qquad (39)$$

Conveniently, this is also a linear first-order ODE with solution derived in the appendix. The following proposition summarizes.

Proposition 4 *Suppose there are potentially two sequential natural experiments, with the same unknown shock arrival rate across periods, $\lambda \in \{\lambda_l, \lambda_h\}$ where $\lambda_h > \lambda_l \geq 0$. Prior to the arrival of shocks, beliefs evolve according to*

$$dB = \frac{(\lambda_h-\lambda_l)B(1-B)dt}{(1-\lambda_l dt)B+(1-\lambda_h dt)(1-B)}.$$

The shadow value of capital prior to the arrival of the first shock is

$$
\begin{aligned}
q_s^1(B) &= \frac{\Pi_s + c}{r+\delta} + \left[\frac{\lambda_l^2[1+\lambda_h/(r+\delta)]C}{(r+\delta+\lambda_l)^2(r+\delta+\lambda_h)}\right] B \\
&\quad + \left[\frac{\lambda_h^2[1+\lambda_l/(r+\delta)]C}{(r+\delta+\lambda_l)(r+\delta+\lambda_h)^2}\right](1-B) \\
&\quad - (\Pi_s + c)\left[\frac{\lambda_l B + \lambda_h(1-B) + \lambda_h\lambda_l/(r+\delta)}{(r+\delta+\lambda_h)(r+\delta+\lambda_l)}\right] \\
c &\equiv \frac{1}{r+\delta}\sum_{\widetilde{s}} \rho^1_{s\widetilde{s}} \Pi_{\widetilde{s}} \\
C &\equiv \sum_{\widetilde{s}} \rho^1_{s\widetilde{s}} \left(\sum_{\widehat{s}} \rho^2_{\widetilde{s}\widehat{s}} (\Pi_{\widehat{s}} - \Pi_{\widetilde{s}})\right).
\end{aligned}
$$

Beliefs just after the arrival of the first shock jump to

$$\widetilde{B}(B) = \frac{\lambda_l B}{B\lambda_l + (1-B)\lambda_h}.$$

The shadow value of capital after the first shock but before the second shock is

$$q^2_{\widetilde{s}}(\widetilde{B}) = \frac{\Pi_{\widetilde{s}}}{r+\delta} + \left[\sum_{\widehat{s}} \rho^2_{\widetilde{s}\widehat{s}} (\Pi_{\widehat{s}} - \Pi_{\widetilde{s}})\right]\left[\frac{\widetilde{B}\lambda_l(r+\delta) + (1-\widetilde{B})\lambda_h(r+\delta) + \lambda_h\lambda_l}{(r+\delta)(r+\delta+\lambda_l)(r+\delta+\lambda_h)}\right].$$

The shadow value of capital after the second and final shock is $q_{\widehat{s}} = \Pi_{\widehat{s}}/(r+\delta)$.

5 Conclusion

This paper considered the problem of interpretation and extrapolation of evidence coming from sequences of seemingly-ideal exogenous policy shocks when the underlying data generating process is not known to either agents or the econometricians studying them. As shown, learning gives rise to "causal parameter drift" even with constant a data generating process. The correct intrepretation of shock responses hinges upon the exact time pattern of realized shocks, as well as (generally unstated) parametric assumptions about priors and potential arrival intensities. Conveniently, closed-form formulae were given for mapping observed shock responses back to theory-implied causal effects (comparative statics).

Appendix: Differential Equation Solutions

Constant Policy Generating Process

We begin by rewriting the ODE as follows:

$$\begin{aligned}
[r+\delta+\lambda(1-B)]q(B) &= \lambda B(1-B)q'(B)+\lambda(1-B)\overline{Q}+\Pi \qquad (40)\\
\overline{Q} &\equiv \sum_{\widetilde{s}\neq s}\rho_{s\widetilde{s}}Q_{\widetilde{s}}.
\end{aligned}$$

This ODE can be written in canonical form as follows:

$$q'(B)-\left[\left(\frac{r+\delta}{\lambda}\right)B^{-1}(1-B)^{-1}+B^{-1}\right]q(B)=\overline{Q}(1-B)^{-1}-\lambda^{-1}[\lambda\overline{Q}+\Pi]B^{-1}(1-B)^{-1}$$

For compactness, this ODE can be expressed as:

$$\begin{aligned}
y'(x)+P(x)y(x) &= \Omega(x) \qquad (41)\\
k &= \frac{r+\delta}{\lambda}\\
m &= \overline{Q}\\
n &= \lambda^{-1}[\lambda\overline{Q}+\Pi]\\
P(x) &= -\left[kx^{-1}(1-x)^{-1}+x^{-1}\right]\\
\Omega(x) &= m(1-x)^{-1}-nx^{-1}(1-x)^{-1}
\end{aligned}$$

We know the solution to this linear first-order ODE is

$$\begin{aligned}
y(x) &= \frac{1}{\mu(x)}\left[\int\mu(x)\Omega(x)dx\right] \qquad (42)\\
\mu(x) &= \exp\left[\int P(x)dx\right].
\end{aligned}$$

Computing the respective indefinite integrals we obtain:

$$\begin{aligned}
\int P(x)dx &= \ln\left[\frac{1}{x}\left(\frac{1-x}{x}\right)^k\right] \qquad (43)\\
&\Rightarrow \mu(x)=\frac{1}{x}\left(\frac{1-x}{x}\right)^k=x^{-k-1}(1-x)^k.
\end{aligned}$$

And the remaining indefinite integral is:

$$\begin{aligned}\int \mu(x)\Omega(x)dx &= \int \frac{1}{x}\left(\frac{1-x}{x}\right)^k m(1-x)^{-1}dx - \int \frac{1}{x}\left(\frac{1-x}{x}\right)^k nx^{-1}(1-x)^{-1}dx \qquad (44)\\ &= -mk^{-1}x^{-k}(1-x)^k + nk^{-1}(k+1)^{-1}x^{-k-1}(1-x)^k(k+x).\end{aligned}$$

Thus we have:

$$\begin{aligned} y &= [x^{k+1}(1-x)^{-k}][-mk^{-1}x^{-k}(1-x)^k] + [x^{k+1}(1-x)^{-k}][nk^{-1}(k+1)^{-1}x^{-k-1}(1-x)^k(k+x)] \\ &= -mxk^{-1} + nk^{-1}(k+1)^{-1}(k+x) \\ &= -mk^{-1}x + n(k+1)^{-1} + nk^{-1}(k+1)^{-1}x \\ &= n(k+1)^{-1} + xk^{-1}[n(k+1)^{-1} - m] \\ &= [\overline{Q} + \lambda^{-1}\Pi]\left(\frac{r+\delta+\lambda}{\lambda}\right)^{-1} + x\left(\frac{\lambda}{r+\delta}\right)\left[[\overline{Q} + \lambda^{-1}\Pi]\left(\frac{r+\delta+\lambda}{\lambda}\right)^{-1} - \overline{Q}\right] \\ &= \frac{\Pi}{r+\delta} + \left(\overline{Q} - \frac{\Pi}{r+\delta}\right)\frac{(1-x)\lambda}{r+\delta+\lambda}.\end{aligned}$$

Heterogeneous Experiments

We begin by rewriting the ODE in canonical form:

$$\begin{aligned} &\frac{\partial}{\partial B}q_s^n(B) - q_s^1(B)\left[\begin{array}{c}\left(\frac{r+\delta}{\lambda_h^n-\lambda_l^n}\right)B^{-1}(1-B)^{-1} + \left(\frac{\lambda_h^n}{\lambda_h^n-\lambda_l^n}\right)B^{-1} \\ +\left(\frac{\lambda_l^n}{\lambda_h^n-\lambda_l^n}\right)(1-B)^{-1}\end{array}\right] \qquad (45)\\ = \;& -\left(\frac{\lambda_l^n}{\lambda_h^n-\lambda_l^n}\sum_{\widetilde{s}}\rho_{s\widetilde{s}}^1 q_{\widetilde{s}}^{n+1}(\beta^{n+1})\right)(1-B)^{-1} \\ & -\left(\frac{\Pi_s}{\lambda_h^n-\lambda_l^n}\right)B^{-1}(1-B)^{-1} \\ & -\left(\frac{\lambda_h^n}{\lambda_h^n-\lambda_l^n}\sum_{\widetilde{s}}\rho_{s\widetilde{s}}^1 q_{\widetilde{s}}^{n+1}(\beta^{n+1})\right)B^{-1}\end{aligned}$$

We write the ODE compactly as:

$$\begin{aligned} &y'(x) - \left[kx^{-1}(1-x)^{-1} + wx^{-1} + z(1-x)^{-1}\right]y(x) \qquad (46)\\ = \;& m(1-x)^{-1} - nx^{-1}(1-x)^{-1} + ox^{-1}\end{aligned}$$

where:

$$\begin{aligned}
k &= \frac{r+\delta}{\lambda_h - \lambda_l} \\
m &= -\left(\frac{\lambda_l}{\lambda_h - \lambda_l}\sum_{\widetilde{s}} \rho^1_{s\widetilde{s}} q^{n+1}_{\widetilde{s}}(\beta^{n+1})\right) \\
n &= \frac{\Pi_s}{\lambda_h - \lambda_l} \\
o &= -\left(\frac{\lambda_h}{\lambda_h - \lambda_l}\sum_{\widetilde{s}} \rho^1_{s\widetilde{s}} q^{n+1}_{\widetilde{s}}(\beta^{n+1})\right) \\
w &= \frac{\lambda_h}{\lambda_h - \lambda_l} \\
z &= \frac{\lambda_l}{\lambda_h - \lambda_l}
\end{aligned}$$

We know the solution to this linear first-order ODE is

$$\begin{aligned}
y(x) &= \frac{1}{\mu(x)}\left[\int \mu(x)\Omega(x)dx\right] \qquad (47) \\
\mu(x) &= \exp\left[\int P(x)dx\right] \\
P(x) &\equiv -\left[kx^{-1}(1-x)^{-1} + wx^{-1} + z(1-x)^{-1}\right] \\
\Omega(x) &\equiv m(1-x)^{-1} - nx^{-1}(1-x)^{-1} + ox^{-1}.
\end{aligned}$$

Computing the indefinite integrals we find:

$$\begin{aligned}
\int P(x)dx &= -k\ln\left(\frac{x}{1-x}\right) - w\ln(x) + z\ln(1-x) \qquad (48) \\
&= k\ln\left[\left(\frac{x}{1-x}\right)^{-1}\right] - w\ln(x) + z\ln(1-x) \\
&= k\ln\left(\frac{1-x}{x}\right) - w\ln(x) + z\ln(1-x) \\
&= \ln\left[\left(\frac{1-x}{x}\right)^{k}\right] - \ln\left(x^{w}\right) + \ln\left[(1-x)^{z}\right] \\
&= \ln\left[x^{-w}\left(\frac{1-x}{x}\right)^{k}(1-x)^{z}\right] \\
&= \ln\left[x^{-k-w}(1-x)^{k+z}\right]
\end{aligned}$$

Thus:

$$\mu(x) = x^{-(k+w)}(1-x)^{k+z}. \tag{49}$$

And we know:

$$\begin{aligned}
k+w &= \frac{r+\delta+\lambda_h}{\lambda_h-\lambda_l} \equiv \kappa_1 \qquad (50)\\
k+z &= \frac{r+\delta+\lambda_l}{\lambda_h-\lambda_l} \equiv \kappa_2\\
&\Rightarrow \ \mu(x) = x^{-\kappa_2-1}(1-x)^{\kappa_2}.
\end{aligned}$$

Computing the remaining indefinite integral we find:

$$\begin{aligned}
\int \mu(x)\Omega(x)dx &= \int [x^{-(k+w)}(1-x)^{k+z}][m(1-x)^{-1} - nx^{-1}(1-x)^{-1} + ox^{-1}]dx \qquad (51)\\
&= m\int [x^{-k-w}(1-x)^{k+z}][(1-x)^{-1}]dx - n\int [x^{-k-w}(1-x)^{k+z}][x^{-1}(1-x)^{-1}]dx\\
&\quad +o\int [x^{-k-w}(1-x)^{k+z}x^{-1}]dx\\
&= m\int (x^{-(k+w)}(1-x)^{k+z-1})dx - n\int (x^{-k-w-1}(1-x)^{k+z-1})dx\\
&\quad +o\int x^{-k-w-1}(1-x)^{k+z}dx\\
&= m\int x^{-\kappa_1}(1-x)^{\kappa_2-1}dx - n\int (x^{-(\kappa_1+1)}(1-x)^{\kappa_2-1})dx\\
&\quad +o\int x^{-k-w-1}(1-x)^{k+z}dx\\
&= m\int x^{-\kappa_2-1}(1-x)^{\kappa_2-1}dx - n\int (x^{-(\kappa_2+2)}(1-x)^{\kappa_2-1})dx\\
&\quad +o\int x^{-k-w-1}(1-x)^{k+z}dx\\
&= m\int x^{-\kappa_2-1}(1-x)^{\kappa_2-1}dx - n\int x^{-\kappa_2-2}(1-x)^{\kappa_2-1}dx\\
&\quad +o\int x^{-\kappa_2-2}(1-x)^{\kappa_2}dx\\
&= m\left(-\frac{1}{x^{\kappa_2}\kappa_2}(1-x)^{\kappa_2}\right) - n\left(-\frac{1}{x^{\kappa_2+1}\kappa_2(\kappa_2+1)}(x+\kappa_2)(1-x)^{\kappa_2}\right)\\
&\quad +o(\kappa_2+1)^{-1}(x-1)(1-x)^{\kappa_2}x^{-\kappa_2-1}\\
&= -m\kappa_2^{-1}x^{-\kappa_2}(1-x)^{\kappa_2} + n\kappa_2^{-1}(\kappa_2+1)^{-1}x^{-\kappa_2-1}(x+\kappa_2)(1-x)^{\kappa_2}\\
&\quad -o(\kappa_2+1)^{-1}(1-x)^{\kappa_2+1}x^{-\kappa_2-1}
\end{aligned}$$

Thus:

$$\begin{aligned}
\frac{1}{\mu}\int \mu Q &= -m\kappa_2^{-1}x + n\kappa_2^{-1}\left(\kappa_2+1\right)^{-1}\left(x+\kappa_2\right) - o\left(\kappa_2+1\right)^{-1}\left(1-x\right) && (52)\\
&= -m\kappa_2^{-1}x + n\left(\kappa_2+1\right)^{-1}\left(x\kappa_2^{-1}+1\right) - o\left(\kappa_2+1\right)^{-1}\left(1-x\right) \\
&= -m\left(\frac{\lambda_h-\lambda_l}{r+\delta+\lambda_l}\right)x + n\left(\frac{\lambda_h-\lambda_l}{r+\delta+\lambda_h}\right)\left(x\kappa_2^{-1}+1\right) \\
&\quad -o\left(\frac{\lambda_h-\lambda_l}{r+\delta+\lambda_h}\right)(1-x) && (53)\\
&= -m\left(\frac{\lambda_h-\lambda_l}{r+\delta+\lambda_l}\right)x + n\left(\frac{\lambda_h-\lambda_l}{r+\delta+\lambda_h}\right)\kappa_2^{-1}x \\
&\quad +n\left(\frac{\lambda_h-\lambda_l}{r+\delta+\lambda_h}\right) - o\left(\frac{\lambda_h-\lambda_l}{r+\delta+\lambda_h}\right)(1-x)\,. && (54)
\end{aligned}$$

Substituting in the definitions of (m,n,o) and rearranging terms we find:

$$\frac{1}{\mu}\int \mu Q = \frac{\Pi_s}{r+\delta} + \left[\sum_{\widetilde{s}} \rho^1_{s\widetilde{s}} q^{n+1}_{\widetilde{s}}(\beta^{n+1}) - \frac{\Pi_s}{r+\delta}\right]\left[\frac{[x\lambda_l+(1-x)\lambda_h](r+\delta)+\lambda_l\lambda_h}{(r+\delta+\lambda_l)(r+\delta+\lambda_h)}\right] \tag{55}$$

Heterogeneous Experiments with Memory

We begin by rewriting the ODE in classical form:

$$\begin{aligned}
&\frac{\partial}{\partial B}q^1_s(B) - q^1_s(B)\left[\begin{array}{c}\left(\frac{r+\delta}{\lambda_h-\lambda_l}\right)B^{-1}(1-B)^{-1} \\ +\left(\frac{\lambda_h}{\lambda_h-\lambda_l}\right)B^{-1} + \left(\frac{\lambda_l}{\lambda_h-\lambda_l}\right)(1-B)^{-1}\end{array}\right] && (56)\\
= \;& -(\lambda_h-\lambda_l)^{-1}B^{-1}(1-B)^{-1}\Pi_s \\
& -(\lambda_h-\lambda_l)^{-1}B^{-1}(1-B)^{-1}[B\lambda_l+(1-B)\lambda_h]\sum_{\widetilde{s}}\rho^1_{s\widetilde{s}}q^2_{\widetilde{s}}[\widetilde{B}(B)].
\end{aligned}$$

Focusing on the final term, we know:

$$\sum_{\widetilde{s}}\rho^1_{s\widetilde{s}}q^2_{\widetilde{s}}[\widetilde{B}(B)] = \underbrace{\frac{1}{r+\delta}\sum_{\widetilde{s}}\rho^1_{s\widetilde{s}}\Pi_{\widetilde{s}}}_{c} + \underbrace{\left[\sum_{\widetilde{s}}\rho^1_{s\widetilde{s}}\left(\sum_{\widehat{s}}\rho^2_{\widetilde{s}\widehat{s}}\left(\Pi_{\widehat{s}}-\Pi_{\widetilde{s}}\right)\right)\right]}_{C} H[\widetilde{B}(B)]. \tag{57}$$

And note, we can write:

$$\begin{aligned}
H[\widetilde{B}(B)] &= \frac{\widetilde{B}\lambda_l}{(r+\delta+\lambda_l)(r+\delta+\lambda_h)} + \frac{(1-\widetilde{B})\lambda_h}{(r+\delta+\lambda_l)(r+\delta+\lambda_h)} && (58)\\
&\quad + \frac{\lambda_h\lambda_l}{(r+\delta)(r+\delta+\lambda_l)(r+\delta+\lambda_h)}.
\end{aligned}$$

From the definition of $\widetilde{B}$ we have

$$\begin{aligned} H[\widetilde{B}(B)] &= \frac{\lambda_l^2 B}{(r+\delta+\lambda_l)(r+\delta+\lambda_h)(B\lambda_l+(1-B)\lambda_h)} \\ &\quad + \frac{(1-B)\lambda_h^2}{(r+\delta+\lambda_l)(r+\delta+\lambda_h)(B\lambda_l+(1-B)\lambda_h)} \\ &\quad + \frac{\lambda_h\lambda_l}{(r+\delta)(r+\delta+\lambda_l)(r+\delta+\lambda_h)}. \end{aligned} \tag{59}$$

And thus we have:

$$\begin{aligned} \sum_{\widetilde{s}} \rho^1_{s\widetilde{s}} q^2_{\widetilde{s}}[\widetilde{B}(B)] &= c + \frac{\lambda_l^2 BC}{(r+\delta+\lambda_l)(r+\delta+\lambda_h)(B\lambda_l+(1-B)\lambda_h)} \\ &\quad + \frac{(1-B)\lambda_h^2 C}{(r+\delta+\lambda_l)(r+\delta+\lambda_h)(B\lambda_l+(1-B)\lambda_h)} \\ &\quad + \frac{\lambda_h\lambda_l C}{(r+\delta)(r+\delta+\lambda_l)(r+\delta+\lambda_h)}. \end{aligned}$$

Thus, the ODE for the shadow value of capital can be written as:

$$\begin{aligned} & \frac{\partial}{\partial B} q^1_s(B)_s - \left[\left(\frac{r+\delta}{\lambda_h-\lambda_l}\right) B^{-1}(1-B)^{-1} + \left(\frac{\lambda_h}{\lambda_h-\lambda_l}\right) B^{-1} + \left(\frac{\lambda_l}{\lambda_h-\lambda_l}\right)(1-B)^{-1}\right] q^1_s(B) \\ = \; & -(\lambda_h-\lambda_l)^{-1}B^{-1}(1-B)^{-1}(\Pi_s+c) \\ & -(\lambda_h-\lambda_l)^{-1}B^{-1}(1-B)^{-1}[B\lambda_l+(1-B)\lambda_h]\frac{\lambda_l^2 BC}{(r+\delta+\lambda_l)(r+\delta+\lambda_h)(B\lambda_l+(1-B)\lambda_h)} \\ & -(\lambda_h-\lambda_l)^{-1}B^{-1}(1-B)^{-1}[B\lambda_l+(1-B)\lambda_h]\frac{(1-B)\lambda_h^2 C}{(r+\delta+\lambda_l)(r+\delta+\lambda_h)(B\lambda_l+(1-B)\lambda_h)} \\ & -(\lambda_h-\lambda_l)^{-1}B^{-1}(1-B)^{-1}[B\lambda_l+(1-B)\lambda_h]\frac{\lambda_h\lambda_l C}{(r+\delta)(r+\delta+\lambda_l)(r+\delta+\lambda_h)}. \end{aligned} \tag{60}$$

Simplifying we obtain:

$$\begin{aligned} & \frac{\partial}{\partial B} q^1_s(B)_s - \left[\left(\frac{r+\delta}{\lambda_h-\lambda_l}\right) B^{-1}(1-B)^{-1} + \left(\frac{\lambda_h}{\lambda_h-\lambda_l}\right) B^{-1} + \left(\frac{\lambda_l}{\lambda_h-\lambda_l}\right)(1-B)^{-1}\right] q^1_s(B) \\ = \; & -\left[\frac{\lambda_l^2[1+\lambda_h/(r+\delta)]C}{(\lambda_h-\lambda_l)(r+\delta+\lambda_l)(r+\delta+\lambda_h)}\right](1-B)^{-1} - \left(\frac{\Pi_s+c}{\lambda_h-\lambda_l}\right)B^{-1}(1-B)^{-1} \\ & -\left[\frac{\lambda_h^2[1+[1+\lambda_l/(r+\delta)]C}{(\lambda_h-\lambda_l)(r+\delta+\lambda_l)(r+\delta+\lambda_h)}\right]B^{-1}. \end{aligned} \tag{61}$$

We write the ODE compactly as:

$$y'(x) - \left[kx^{-1}(1-x)^{-1} + wx^{-1} + z(1-x)^{-1}\right] y(x) = m(1-x)^{-1} - nx^{-1}(1-x)^{-1} + ox^{-1} \tag{62}$$

where:

$$\begin{aligned}
k &= \frac{r+\delta}{\lambda_h-\lambda_l} \\
m &= -\left[\frac{\lambda_l^2[1+\lambda_h/(r+\delta)]C}{(\lambda_h-\lambda_l)(r+\delta+\lambda_l)(r+\delta+\lambda_h)}\right] \\
n &= \left(\frac{\Pi_s+c}{\lambda_h-\lambda_l}\right) \\
o &= -\left[\frac{\lambda_h^2[1+[1+\lambda_l/(r+\delta)]C}{(\lambda_h-\lambda_l)(r+\delta+\lambda_l)(r+\delta+\lambda_h)}\right] \\
w &= \frac{\lambda_h}{\lambda_h-\lambda_l} \\
z &= \frac{\lambda_l}{\lambda_h-\lambda_l}
\end{aligned}$$

We know the solution to this linear first-order ODE is

$$\begin{aligned}
y(x) &= \frac{1}{\mu(x)}\left[\int \mu(x)\Omega(x)dx\right] \qquad (63) \\
\mu(x) &= \exp\left[\int P(x)dx\right] \\
P(x) &\equiv -\left[kx^{-1}(1-x)^{-1}+wx^{-1}+z(1-x)^{-1}\right] \\
\Omega(x) &\equiv m(1-x)^{-1}-nx^{-1}(1-x)^{-1}+ox^{-1}.
\end{aligned}$$

These are the same indefinite integrals as in the preceding analysis of Heterogeneous Experiments without memory. From the calculations of those integrals we know (see equation (52)):

$$\begin{aligned}
\frac{1}{\mu}\int \mu Q &= -m\left(\frac{\lambda_h-\lambda_l}{r+\delta+\lambda_l}\right)x+n\left(\frac{\lambda_h-\lambda_l}{r+\delta+\lambda_h}\right)\kappa_2^{-1}x \qquad (64) \\
&\quad +n\left(\frac{\lambda_h-\lambda_l}{r+\delta+\lambda_h}\right)-o\left(\frac{\lambda_h-\lambda_l}{r+\delta+\lambda_h}\right)(1-x)\,.
\end{aligned}$$

Using the definitions of (m, n, o) we find:

$$\begin{aligned}
\frac{1}{\mu}\int \mu Q &= \left[\frac{\lambda_l^2[1+\lambda_h/(r+\delta)]C}{(\lambda_h-\lambda_l)(r+\delta+\lambda_l)(r+\delta+\lambda_h)}\right]\left(\frac{\lambda_h-\lambda_l}{r+\delta+\lambda_l}\right)x \\
&\quad +\left(\frac{\Pi_s+c}{\lambda_h-\lambda_l}\right)\left(\frac{\lambda_h-\lambda_l}{r+\delta+\lambda_h}\right)\frac{(\lambda_h-\lambda_l)}{(r+\delta+\lambda_l)}x \\
&\quad +\left(\frac{\Pi_s+c}{\lambda_h-\lambda_l}\right)\left(\frac{\lambda_h-\lambda_l}{r+\delta+\lambda_h}\right) \\
&\quad +\left[\frac{\lambda_h^2[1+[1+\lambda_l/(r+\delta)]C}{(\lambda_h-\lambda_l)(r+\delta+\lambda_l)(r+\delta+\lambda_h)}\right]\left(\frac{\lambda_h-\lambda_l}{r+\delta+\lambda_h}\right)(1-x)\,.
\end{aligned}$$

Grouping terms we obtain:

$$
\begin{aligned}
\frac{1}{\mu}\int \mu Q &= \left(\frac{\Pi_s + c}{r+\delta}\right) + \left[\frac{\lambda_l^2[1+\lambda_h/(r+\delta)]C}{(r+\delta+\lambda_l)^2(r+\delta+\lambda_h)}\right]x + \left[\frac{\lambda_h^2[1+\lambda_l/(r+\delta)]C}{(r+\delta+\lambda_l)(r+\delta+\lambda_h)^2}\right](1-x) \\
&+ \left(\frac{(\lambda_h-\lambda_l)(\Pi_s+c)}{(r+\delta+\lambda_h)(r+\delta+\lambda_l)}\right)x + (\Pi_s+c)\left(\frac{1}{r+\delta+\lambda_h} - \frac{1}{r+\delta}\right) \\
&= \left(\frac{\Pi_s + c}{r+\delta}\right) + \left[\frac{\lambda_l^2[1+\lambda_h/(r+\delta)]C}{(r+\delta+\lambda_l)^2(r+\delta+\lambda_h)}\right]x + \left[\frac{\lambda_h^2[1+\lambda_l/(r+\delta)]C}{(r+\delta+\lambda_l)(r+\delta+\lambda_h)^2}\right](1-x) \\
&+ \left(\frac{(\lambda_h-\lambda_l)(\Pi_s+c)}{(r+\delta+\lambda_h)(r+\delta+\lambda_l)}\right)x - (\Pi_s+c)\left(\frac{\lambda_h}{(r+\delta)(r+\delta+\lambda_h)}\right) \\
&= \left(\frac{\Pi_s + c}{r+\delta}\right) + \left[\frac{\lambda_l^2[1+\lambda_h/(r+\delta)]C}{(r+\delta+\lambda_l)^2(r+\delta+\lambda_h)}\right]x + \left[\frac{\lambda_h^2[1+\lambda_l/(r+\delta)]C}{(r+\delta+\lambda_l)(r+\delta+\lambda_h)^2}\right](1-x) \\
&+ (\Pi_s+c)\left[\frac{\lambda_h x}{(r+\delta+\lambda_h)(r+\delta+\lambda_l)} - \frac{\lambda_l x}{(r+\delta+\lambda_h)(r+\delta+\lambda_l)} - \frac{\lambda_h}{(r+\delta)(r+\delta+\lambda_h)}\right] \\
&= \left(\frac{\Pi_s + c}{r+\delta}\right) + \left[\frac{\lambda_l^2[1+\lambda_h/(r+\delta)]C}{(r+\delta+\lambda_l)^2(r+\delta+\lambda_h)}\right]x + \left[\frac{\lambda_h^2[1+\lambda_l/(r+\delta)]C}{(r+\delta+\lambda_l)(r+\delta+\lambda_h)^2}\right](1-x) \\
&- (\Pi_s+c)\left[\frac{\lambda_l x + \lambda_h(1-x) + \lambda_h\lambda_l/(r+\delta)}{(r+\delta+\lambda_h)(r+\delta+\lambda_l)}\right]
\end{aligned}
$$