Title

Learning, parameter drift, and the credibility revolution

Authors

Hennessy, Christopher A
Livdan, Dmitry

Publication Date

2021

DOI

10.1016/j.jmoneco.2020.02.003

Peer reviewed

Learning, Parameter Drift, and the Credibility Revolution[⋆]

Christopher A. Hennessy[a,*], Dmitry Livdan[b]

[a]*London Business School, Regent's Park, London, NW1 4SA, U.K.*
[b]*Walter A. Haas School of Business, University of California, Berkeley, CA 94720, USA*

Abstract

This paper analyses extrapolation and inference using tax experiments in dynamic economies when shock processes are latent regime-shifting Markov chains. Belief revisions result in severe parameter drift: Response signs and magnitudes vary widely over time despite ideal exogeneity. Even with linear causal effects, shock responses are non-linear, preventing direct extrapolation. Analytical formulae are derived for extrapolating responses or inferring causal parameters. Extrapolation and inference hinges upon shock histories and correct assumptions regarding potential data generating processes. A martingale condition is necessary and sufficient for shock responses to directly recover comparative statics, but stochastic monotonicity is insufficient for correct sign inference.

Keywords: Natural Experiment, Causality, Uncertainty, Learning.
JEL: E62, E63, G18, G28, G38, H00

[⋆]We thank seminar participants at Bocconi, Banca di Italia, and LBS for feedback. Thanks also to Andrew Abel, Jessica Wachter, and Howard Kung for helpful feedback. Finally, we are grateful to Urban Jermann (the editor), Francesco Bianchi (the associate editor), and an anonymous referee for substantive recommendations that allowed us to generalize the model.

[*]CEPR and ECGI. Corresponding author

> The major contributions of twentieth century econometrics to knowledge were the definition of causal parameters when agents are constrained by resources and markets and causes are interrelated, the analysis of what is required to recover causal parameters from data (the identification problem), and clarification of the role of causal parameters in policy evaluation and in forecasting the effects of policies never previously experienced.
> –James Heckman (2000)

1. Introduction

Angrist and Pischke (2010) argue that exploitation of quasi-natural experiments amounts to a "credibility revolution" in resolving the causal parameter identification problem. They go on to criticize macroeconomists for failing to share their revolutionary zeal, arguing that "today's macro agenda is empirically impoverished... The theory-centric macro fortress appears increasingly hard to defend."

Notwithstanding the principled objections of Sims (2010), Keane (2010) and Rust (2010), amongst others, a fair reading of the state of play is that the model-light empirical methodology recommended by Angrist and Pischke (2010) is presently in the ascendancy. This view also appears to have gained ground with some macroeconomists. For example, Romer (2016) questions identification strategies in macroeconomics, while Narayana Kocherlakota (2018) argues "there has been a revolution in applied microeconometrics in the use of atheoretical statistical methods... a similar change could be of value in applied macroeconomics." Romer and Romer (2014) argue, "In microeconomic settings, it is often possible to identify natural experiments where it is clear that differences among economic actors are not the result of confounding factors."

In part, the appeal of Angrist and Pischke's recommended methodological tool-kit is the heuristic connection between "experiments" and "causal effects." Apparently, many consider it to be *a priori* obvious that quasi-natural experiments recover causal effects if exploited shocks can be shown to be exogenous. This accounts for the narrow focus of many econometricians on finding sources of exogenous variation, with little attention devoted to mapping coefficients back to causal parameters. This view is the hallmark of the influential textbook of Angrist and Pischke (2009), *Mostly Harmless Econometrics: An Empiricist's Companion*. They write, "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." They maintain, "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."

Undermining such assertions of credibility, Angrist and Pischke (2009, 2010) never formally demonstrate the connection between quasi-natural experiments and causal parameters. To the contrary, Hennessy and Strebulaev (2019) show that in dynamic economies, responses to exogenous shocks generally fail to recover two important causal parameters: theory-implied causal effects (comparative statics) and policy-invariant adjustment cost parameters determining causal effect magnitudes. However, responses to specific policy variable transitions do forecast responses to identical policy variable transitions in the setting they consider.

In fact, there is a more obvious observation casting doubt on assertions of inherent credibility of natural experiments: If an empirical methodology is credible, those applying the methodology should arrive at similar quantitative estimates regarding the magnitude of causal parameters. However, the stock of widely conflicting quantitative evidence being accumulated in fields such as labor, development, environmental, and public economics suggests the presence of *parameter drift*, or time-varying econometric estimates of quantities that are, by definition, constant over time. For example, contrary to Hennessy and Strebulaev (2019), historical shock responses do not even appear to be good forecasters of future shock responses.

As shown by Lucas (1976), whose focus was on parameters underpinning large-scale macroeconometric models, a potential source of parameter drift is a change in the underlying stochastic process–and this is true if experiment shock response magnitudes are treated as the causal parameter of interest. Conveniently, progress has been made in developing quasi-structural methods for recovering causal parameters in quasi-experimental settings featuring dynamic uncertainty and/or changes in underlying stochastic processes, e.g. Heckman and Navarro (2007) and Hennessy and Strebulaev (2019). However, reduced-form econometricians often object to using these methods since they demand making "strong" distributional assumptions. In turn, reluctance to make distributional assumptions reflects the fact that applied econometricians are often uncertain about the data generating process for the shocks they exploit. In fact, this type of model uncertainty is often invoked as a defense amongst those recommending reduced-form quasi-experimental methods over structural estimation.

It must be conceded that in many applied settings econometricians and the agents they study are unlikely to be certain of the true underlying process generating the (exogenous) shocks being exploited. But what implications does this type of model uncertainty have for quasi-experimental inference, and what can be done about it? The objective of this paper is to address these questions, and clarify the issues, using a transparent *analytical* framework. To do so, we follow the rational expectations approach of Hansen and Sargent (2010) in treating agents and econometricians symmetrically. In particular, we give the reduced-form econometrician the argument that there is uncertainty regarding the underlying stochastic process generating the exogenous shocks being exploited in the pursuit of causal parameters. But then, imposing the symmetry demanded by rational expectations, we assume that the agents being observed by the econometrician also do not know the underlying shock generating process. Rather, agents and econometricians know the set of potential models and engage in Bayesian updating. Within this context, we derive *closed-form* expressions clarifying the relationship between evidence from natural experiments and causal effect parameters.

We consider the following economic setting. An econometrician seeks to empirically estimate causal effect parameters as implied by a canonical dynamic theory: investment by firms using a linear-quadratic technology. To fix ideas, we focus on linear tax rate shocks that reduce the return to investment and analyze their causal impact, although our analysis applies to any linear profit shock. Importantly, as shown, the linear-quadratic technology gives rise to the classical linear causal effect econometric framework. In the linear causal effect framework, changes in the dependent variable (here investment) are linear in changes to the independent variable (here tax rates). The causal effect parameter to be estimated by the econometrician can be a time-homogeneous comparative static, a policy-invariant technological parameter, or a shock response forecast.

The econometrician exploits tax rate shocks that are "ideal" in the Angrist-Pischke sense that endogeneity and selection are not a concern. In particular, the tax rate is governed by an independent N-state continuous-time Markov chain with regime shifting. All agents, including the econometrician, face model uncertainty. We consider a very general form of model uncertainty: agents may be uncertain about tax shock arrival probabilities and/or the probability distribution governing tax rate transitions.[1] Formally, we consider that the instantaneous Markov transition matrix can assume one of J potential values, with instantaneous switches across matrices possible. Firms are embedded in a general equilibrium setting where the marginal product of capital is proportional to exogenous aggregate output.

The most important negative findings are as follows. First, uncertainty about the underlying stochastic process severely complicates the mapping between observed shock responses and causal parameters. For example, correct interpretation hinges upon correctly stipulating the set of potential data generating processes, correctly stipulating the probability weights placed on the alternative processes before the shock, and correctly stipulating how beliefs will change after a given shock. This contradicts Angrist and Pischke's (2009) bold assertion that natural experiments have a "simple interpretation" and also serves as a counterweight to the conventional wisdom that model uncertainty somehow tilts the balance in favor of reduced-form inference. Natural experiments only have a simple interpretation if one takes them at face value. Once one uses a parable economy to mimic such experiments, as we do, it becomes apparent that making valid inferences requires making assumptions about functional forms and data generating processes, just as structural work requires. Moreover, model uncertainty, specifically uncertainty about underlying data generating processes, confounds inference in natural experiments in much the same manner as structural work. The only distinction is that structural work puts these issues into the open while quasi-experimental work maintains they are not an issue, until objections are raised, at which point it is argued that the assumptions are implicit yet somehow absent from the textbooks.

Second, if the underlying stochastic process is latent, causal parameter drift will be commonplace in shock-based inference. Simply put, there is no *a priori* reason to expect econometricians estimating shock responses at different points in time to produce similar estimates, even if the shocks are identical. Phrased differently, with learning, past shock responses are poor unconditional forecasters of future shock responses.

[1] An early version of this paper considered only two possible shock intensities. We thank the editors and referee for suggesting this extension.

Intuitively, endogenous time-variation in beliefs gives rise to time-variation in shock responses. Importantly, this is so even if we assume the true data generating process is known to be constant, so that the Lucas critique does not apply.

Third, it is shown that shock responses do not necessarily recover the correct sign of the theory-implied causal effect. That is, the problem of causal parameter drift is not confined to magnitudes but extends also to signs. Intuitively, without context, a tax rate cut appears to be good news. However, the specific tax cut may not be viewed as good news by Bayesian agents. After all, they might have expected a larger cut. Or the specific tax cut may cause them to expect less generous tax cuts in the future. As a practical matter, such results call into doubt the interpretation and utilization of elasticity estimates shaping policy. 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." What Slemrod fails to account for is the fact that the information content of shocks varies systematically with waiting times, with more evidence often being worse evidence.

Fourth, an important mechanism made clear within our framework is that shock responses hinge not only on the beliefs held by agents just prior to the shock arriving, but depend also on the belief revision that a given natural policy experiment brings about. As we show, this belief revision effect can radically change both the sign and magnitude of shock responses. For example, firms may respond to a tax rate cut by cutting their investment if it causes them to place lower weight on relatively favorable data generating processes.

Fifth, although we consider a setting in which causal effects are linear in the size of tax rate changes, there is no reason to assume that shock responses are symmetrical or proportional to shock sizes. This calls into question the common practice of extrapolating shock responses based upon size. Simply put, even with a technology consistent with linear theory-implied causal effects, shock responses are not generally linear. Intuitively, there is no *a priori* reason to assume that belief revisions are symmetrical or proportional, and belief revisions are fundamental in the decomposition of shock responses.

Finally, we extend the model to allow for aggregate uncertainty. Specifically, we follow Veronesi (2000) in assuming the instantaneous drift rate of aggregate output follows a latent regime shifting process. As shown, such macroeconomic uncertainty further complicates the mapping between shock responses and causal effects. In particular, the correct interpretation of natural experiments hinges upon correctly specifying beliefs about the underlying data generating processes driving *both* microeconomic and macroeconomic shocks. In this sense, applied microeconometricians must confront many of the same issues confronting macroeconometricians, even if the tool-kits differ.

The constructive contribution of the paper is to illustrate how to account for learning and dynamic model uncertainty in shock-based inference, so that the problem of causal parameter drift can be addressed operationally. We first provide analytical expressions for mapping observed shock responses to causal effect parameters, specifically, comparative statics, policy-invariant technological parameters, or shock response forecasts. Essentially, the econometrician must impose upon herself the "communism of models" of Sargent (2005) with empirically observed shock responses being adjusted using the same real-time information set, and beliefs, as the agents being studied. With consistent belief adjustments, shock responses measured at different points can be rendered comparable and/or converted back to comparative statics. Further, unbiased estimates of deep technological parameters can be extracted from shock responses.

As a second constructive result, we derive an auxiliary identifying assumption, beyond random assignment, that is necessary and sufficient for shock responses to directly recover theory-implied causal effects (comparative statics) in economies where agents and econometricians learn over time: For all potential data generating processes the tax rate is a martingale. Intuitively, Hennessy and Strebulaev (2019) show that in economies where profitability is driven by a *known* Markov chain, martingale profitability is sufficient for shadow values to behave as if shocks are completely unanticipated and permanent, so that shock responses directly recover comparative statics. In this paper, we show an analogous result obtains even if agents do not know the data generating process. However, in contrast to Hennessy and Strebulaev (2019), we show that stochastic monotonicity of all potential data generating processes is insufficient to ensure shock responses correctly recover the sign of theory-implied causal effects.

The present paper shares with Gomes (2001) and Moyen (2004) the idea of using a canonical neoclassical model to shed light on empirical evidence. Their analysis is numerical and they do not analyze natural experiments or learning. The linear-quadratic stock accumulation model used in the paper follows Abel and

Eberly (1994) and Abel and Eberly (1997), but incorporates learning. Jovanovic (1982) analyzes the effect of learning on firm dynamics. Learning has featured in subsequent analysis of investment decisions by Alti (2003), Decamps and Mariotti (2004), and Bouvard (2014).

Our framework can be seen as straddling two strands of the macro-finance literature on learning. One strand, exemplified by Bianchi and Melosi (2016), seeks to incorporate learning dynamics within rich Markov-switching DSGE settings in a computationally tractable way amenable to estimation, as in Bianchi and Melosi (2019). Another strand of the literature, exemplified by Veronesi (2000), considers simpler environments admitting analytical solutions. Although we allow for a richer learning environment than Veronesi, we still pursue and obtain analytical solutions. This objective arises from our view that it is unlikely to expect reduced-form empiricists to embrace numerical/structural methods. Moreover, analytical solutions lay bare the key mechanisms to audiences prone to labeling numerical solutions as a " black box." Of course, none of the learning papers discussed analyzes implications for empirical work exploiting natural experiments. In contrast, Hennessy and Strebulaev (2019) do analyze natural experiments, but they do not allow for the possibility of model uncertainty.

The present paper shares with Keane and Wolpin (2002) the notion that one must account for dynamics and randomness in order to correctly infer causal effects. However, there are numerous important differences. First, they analyze a granular dynamic model of contraceptive use and welfare participation. We offer a more general/abstract analysis of the effect of dynamics and uncertainty on shadow values, the key determinant of optimal accumulation of stock variables. Second, they offer numerical solutions featuring polynomial approximations while we present closed-form solutions amenable to direct analysis and back-of-the-envelope adjustments. Finally, and most importantly, we consider the problem of causal inference in economies in which agents do not know the underlying stochastic process.

The remainder of the paper is organized as follows. Section 2 describes the baseline economic setting. Section 3 presents characterization of optimal investment and shock responses under microeconomic uncertainty. Section 4 illustrates the potential quantitative significance of parameter drift in natural experiments using the realized time-series of historical changes in effective corporate income tax rates. Section 5 extends the baseline model to incorporate macroeconomic uncertainty. Section 6 concludes.

2. Baseline Economic Setting

We consider a general equilibrium (GE) setting that is sufficiently tractable analytically to admit closed-form solutions, even as we consider general forms of microeconomic and macroeconomic uncertainty. This section describes the baseline economic setting. In this baseline setting, the stochastic process for aggregate output is common knowledge, with uncertainty being confined to the nature of tax rate shocks that are "microeconomic" in the sense of leaving aggregate output unchanged.

2.1. Technology

Time is continuous and the horizon is infinite. Uncertainty is modeled by a complete probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The only resource is divisible land. The total amount of land is $\overline{K}$, where $\overline{K}$ is an arbitrarily large constant. The land is uniformly covered with Lucas trees. Each unit of land provides an instantaneous flow of the perishable consumption good (fruit) $X_t dt$. The output process X is a geometric Brownian motion which evolves under the physical measure $\mathbb{P}$ as follows:

$$\begin{aligned} dX_t &= \mu X_t dt + \sigma dW^P \\ X_0 &> 0. \end{aligned} \tag{1}$$

Each parcel of land is owned by either the government or corporations. Regardless of who owns a parcel of land, its respective fruit can be harvested at zero cost. The corporate sector consists of a measure-one continuum of identical non-cooperative firms. Aggregate corporate land at time t is K_t and aggregate corporate revenue is $K_t X_t dt$. The government stands ready to buy and sell $I_t dt$ units of land in exchange for a land fee $(I_t + \gamma I_t^2)dt$. The government levies a tax at rate $T_t \in [0, 1)$ on corporate revenue, implying corporate tax proceeds $T_t K_t X_t dt$. The government redistributes in lump sum fashion corporate taxes, land fees, and fruit harvested on government land. By construction, the posited technology fixes aggregate output at $\overline{K} X_t dt$.

The economy has a representative agent with power-function utility. In order for markets to clear, the representative agent must find it optimal to consume aggregate output. As is well-known, the risk-free rate (r) and risk-premium (θ) in such an economy are constants, and any asset can be priced by discounting at rate r expected cash flow under the risk-neutral measure $\mathbb{Q}$.[2] The dynamics of the output process under the risk-neutral measure are given by

$$dX_t = (\mu - \sigma\theta)X_t dt + \sigma dW^Q. \tag{2}$$

A corporation's instantaneous investment $(I_t)_{t\geq 0}$ must be right-continuous and progressively measurable with respect the augmented filtration generated by X and T. To maintain consistency with the investment literature, which generally analyzes investment in depreciating capital goods, assume that at each instant the government seizes from each corporation a fraction δ of its land holdings. The implied law of motion for corporate sector land is

$$dK_t = (I_t - \delta K_t)dt. \tag{3}$$

The tax rate can take one of $N \geq 2$ values. In tax state S the tax rate is T_S. Of course, the tax rate/state are common knowledge. The tax rate T evolves a continuous-time Markov chain. At any instant, the Markov chain can driven by one of $J \geq 2$ transition matrices, with matrices indexed by i or j below. The true instantaneous Markov matrix is not observed by any agent. Supposing we are in tax state S, then if j were in fact the true instantaneous Markov matrix, then over the next infinitesimal time interval dt there is probability $\lambda_S^j dt$ that a new tax rate state S' will be chosen according to the distribution function $\rho_{SS'}^j$. Notice, the law of motion for the tax rate varies with the true underlying Markov matrix *and* the current tax state.

Given true initial Markov matrix j, over the next infinitesimal time interval dt there is probability $\phi_j dt$ of a transition to a new matrix according to the probability the distribution function π_{ji}. Notice this setup allows for uncertainty regarding shock probabilities and/or shock distribution functions, and allows for both constant and regime shifting data generating processes.

By construction we rule out endogeneity/selection bias by assuming T and X are independent stochastic processes. For brevity, we summarize this important assumption as:

$$T \perp X. \tag{4}$$

Of course, applied microeconometricians devote great attention to addressing concerns arising from endogeneity. Our objective is to strip away this concern in order to show that establishing independence of shocks is a far cry from establishing identification of causal effects.

2.2. The Econometrician

We suppose now that there is a "real-world" applied microeconometrician who performs shock-based causal inference within this economy. To begin, we must formally define the objects this econometrician would like to infer.

The traditional definition of a causal effect is a comparative static. Heckman (2000) writes, "Comparative statics exercises formalize Marshall's notion of a ceteris paribus change which is what economists mean by a causal effect." Athey, Milgrom and Roberts (1998) write, " most of the testable implications of economic theory are comparative static predictions." Analytical comparative statics generally contemplate infinitesimal changes in causal variables. Numerical comparative statics contemplate discrete changes in causal variables. Problematically, Angrist and Pischke (2009) never formally define the theoretical objects natural experiments recover. Nevertheless, their textbook implies that natural experiments recover objects most similar to numerical comparative statics. They write, " 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." Of course, quantitative theorists make counterfactual predictions by simulating parable economies under alternative assumptions regarding causal parameters.

In our parable economy, the *theory-implied causal effect* (CE) is the comparative static of investment with respect to T. With the tax rate treated as a parameter permanently fixed at T, rather than as a stochastic

[2]See Goldstein, Ju and Leland (2001) for example.

process, the shadow value of a unit of land is

$$Q_t = \frac{(1-T)X_t}{r+\delta-\mu+\sigma\theta}. \tag{5}$$

The optimal instantaneous control policy in such a constant tax rate economy, call it I_t^{**}, entails investing up to the point that the shadow value of land is just equal to marginal costs:

$$Q_t = 1 + 2\gamma I_t^{**} \Rightarrow I_t^{**} = \left(\frac{1}{2\gamma}\right)\left[\left(\frac{1-T}{r+\delta-\mu+\sigma\theta}\right)X_t - 1\right]. \tag{6}$$

From the preceding two equations we obtain the following theory-implied causal effects, respectively, for infinitesimal changes and discrete changes in the corporate tax rate from T_S to $T_{S'}$:

$$\begin{aligned} CE &\equiv \frac{\partial I^{**}}{\partial T} = -\left(\frac{1}{2\gamma}\right)\left(\frac{1}{r+\delta-\mu+\sigma\theta}\right)X_t \\ CE_{SS'} &\equiv I_{S'}^{**} - I_S^{**} = \left(\frac{1}{2\gamma}\right)\left(\frac{1}{r+\delta-\mu+\sigma\theta}\right)X_t \times (T_S - T_{S'}). \end{aligned} \tag{7}$$

Notice, the posited linear-quadratic technology gives rise to the classical linear causal effects econometric model. In particular, the theory-implied causal effect is proportional to the size of the change in the causal variable T.

In many cases researchers are interested in directly estimating policy-invariant structural parameters. For example, Summers (1981) attempts to infer the investment cost parameter γ based upon regressions of investment rates on Tobin's Q. In this paper, we consider that the econometrician wants to instead exploit responses to "clean" tax rate shocks in order to infer γ. Alternatively, we consider that the econometrician may want to predict future shock responses based upon an observed shock response. That is, the econometrician may want to extrapolate past shock responses into future shock responses.

3. Microeconomic Model

This section presents an analytical characterization of optimal investment and shock responses under "microeconomic uncertainty," which is uncertainty that does not relate to aggregate output.

3.1. Preliminaries: No Uncertainty

To motivate the solution with uncertainty, it is useful to consider first firm behavior absent uncertainty. In particular, consider an investment program indexed by j, with j representing a known data generating process. The Hamilton-Jacobi-Bellman (HJB) equation is:

$$\begin{aligned} rV^j(K,X,S) &= \max_I \quad V_k^j(I-\delta K) + V_x^j(\mu-\sigma\theta)X + \frac{1}{2}\sigma^2 X^2 V_{xx}^j \\ &\quad +\lambda_S^j \sum_{S'\neq S} \rho_{SS'}^j [V^j(K,X,S') - V^j(K,X,S)] + (1-T_S)KX - I - \gamma I^2. \end{aligned} \tag{8}$$

The HJB equation is an equilibrium condition demanding that the risk-neutral expecting holding return on the firm's stock is just equal to the risk-free rate. As shown above, the holding return consists of capital gains due to infinitesimal changes in the diffusion processes, plus discrete capital gains due to changes in the tax rate, plus dividends.

As shown by Abel and Eberly (1997), with benefits that are linear in the stock and adjustment costs that are independent of the stock, the value function takes the separable form:

$$V^j(K,X,S) = KQ^j(X,S) + G^j(X,S). \tag{9}$$

In fact, separability of the value function between assets in place and growth options will continue to hold even as we incorporate learning. As we show, separability is verified as HJB equation decouples into two

PDEs, with only one of the PDEs involving K, with K entering as a scalar in fact. This K-scaled PDE pins down Q. In fact, this same argument is employed by Abel and Eberly (1997).

Isolating those terms in the HJB equation involving the investment policy I, the optimal instantaneous investment solves:

$$\begin{aligned} &\max_I \quad Q^j(X,S)I - I - \gamma I^2 \qquad (10)\\ \Rightarrow \quad & I_S^* = \frac{Q^j(X,S)-1}{2\gamma};\ S=1,...,N \\ \Rightarrow \quad & I_S^* Q(X,B,S) - I_S^* - \gamma I_S^{*2} = \frac{[Q^j(X,S)-1]^2}{4\gamma} \end{aligned}$$

Since the HJB equation must hold point-wise, the terms scaled by K must equate. It follows that the shadow value of capital must satisfy:

$$(r+\delta+\lambda_S^j)Q^j(X,S) = (\mu-\sigma\theta)XQ_x^j(X,S) + \frac{1}{2}\sigma^2X^2Q_{xx}^j(X,S) + \lambda_S^j \sum_{S'\neq S} \rho_{SS'}^j Q^j(X,S') + (1-T_S)X. \quad (11)$$

We conjecture the shadow value is linear in X and thus write:

$$Q^j(X,S) = X\Psi_S^j$$

where $\mathbf{\Psi}^j$ is an N dimensional vector of constants to be determined. Substituting the preceding expression into the shadow value equation we obtain the following condition:

$$(r+\delta-\mu+\sigma\theta+\lambda_S^j)\Psi_S^j = \lambda_S^j \sum_{S'\neq S} \rho_{SS'}^j \Psi_{S'}^j + (1-T_S). \quad (12)$$

From the preceding equation it follows that the vector of shadow value constants $\mathbf{\Psi}^j$ solves a linear system. We thus have the following proposition.

Proposition 1. *If there is no model uncertainty and the tax rate evolves according to a known continuous-time Markov chain j, then the tax-state-contingent shadow value of capital is*

$$\widetilde{\mathbf{Q}}(X) = X\widetilde{\mathbf{\Psi}}^j$$

where the N state-contingent shadow value constants $\{\widetilde{\Psi}_S^j\}$ solve the following system of linear equations

$$\begin{aligned} 1-T_1 &= (r+\delta-\mu+\sigma\theta+\lambda_1^j)\widetilde{\Psi}_1^j - \lambda_1^j \sum_{S'\neq 1} \rho_{1S'}^j \widetilde{\Psi}_{S'}^j. \\ &... \\ 1-T_N &= (r+\delta-\mu+\sigma\theta+\lambda_N^j)\widetilde{\Psi}_N^j - \lambda_N^j \sum_{S'\neq N} \rho_{NS'}^j \widetilde{\Psi}_{S'}^j. \end{aligned}$$

Hennessy and Strebulaev (2019) derive a similar expression for shadow values under a known stochastic process albeit in a simpler partial equilibrium setting without the geometric Brownian motion X capturing aggregate risk. Before closing this subsection, we anticipate that in certain cases, shadow values under model uncertainty will represent belief weighted averages of the preceding shadow values absent uncertainty. As in the proposition, tildes will be used to represent shadow values and shadow value constants absent model uncertainty.

3.2. Shadow Values under Uncertainty

Suppose now that agents do not know the tax generating process. To begin, let $\mathbf{B}$ denote a vector of dimension J representing agents' probability assessments regarding the current instantaneous Markov matrix. Consider first an instant dt over which no tax rate change occurs. Applying Bayes' law we have:

$$B_j + dB_j = \frac{B_j(1-\phi_j dt)(1-\lambda_S^j dt) + \sum_{i\neq j} B_i \phi_i \pi_{ij} dt (1-\lambda_S^i dt)}{1-\sum_i B_i \lambda_S^i dt} \tag{13}$$

$$\Rightarrow \quad dB_j = \frac{\left[B_j\left(\sum_i B_i\lambda_S^i - \lambda_S^j\right) + \sum_{i\neq j} B_i\phi_i\pi_{ij} - B_j\phi_j\right] dt}{1 - dt\sum_i B_i\lambda_S^i}.$$

The intuition for the preceding equation is as follows. First, if there were no possibility of a switch in the underlying Markov matrix, then B_j would increase in response to no tax rate change if λ_S^j were to fall below the expected value of λ_S given beliefs the preceding instant. This effect is captured by the first term in the numerator of the second equation. The last two terms in the numerator capture changes in beliefs due to expected transitions into and out of Markov matrix j. As another special case of this law of motion, note that if there were no possibility of switches across Markov matrices, and if the shock arrival rate were equal across all j, then beliefs would be constant over time intervals with no tax rate change.

Consider next the evolution of beliefs in the event of a transition from tax state S to state S'. Applying Bayes' rule and dropping terms smaller than infinitesimal dt, we find that after a tax rate change beliefs will generally exhibit a discrete jump to[3]

$$\widetilde{B}_j(\mathbf{B}) = B_j \times \frac{\lambda_S^j \rho_{SS'}^j}{\sum_i B_i \lambda_S^i \rho_{SS'}^i}. \tag{14}$$

The preceding equation shows that after a tax rate change, the probability weight placed on Markov matrix j will increase if it features a higher instantaneous probability of a jump from S to S' relative to the expected probability of such a jump given beliefs the preceding instant. Of course, this is a central point of our paper: the arrival of an experiment itself can be responsible for large revisions of beliefs. And, as shown below, such belief revisions can severely cloud causal inference, and even bring about sign reversals.

In the interest of brevity we present here key steps in the characterization of investment and shadow values. All intermediate steps can be found in the Online Appendix. The HJB equation is:

$$\begin{aligned} & rV(K,X,\mathbf{B},S)dt \qquad (15) \\ = \; & \max_I \quad \left[V_k(I-\delta K)dt + V_x(\mu-\sigma\theta)Xdt + \frac{1}{2}\sigma^2 X^2 V_{xx} dt\right]\left[1 - dt\sum_i B_i\lambda_S^i\right] \\ & + \sum_j V_{b_j}\left(\frac{\left[B_j\left(\sum_i B_i\lambda_S^i - \lambda_S^j\right) + \sum_{i\neq j} B_i\phi_i\pi_{ij} - B_j\phi_j\right] dt}{1-dt\sum_i B_i\lambda_S^i}\right)\left(1 - dt\sum_i B_i\lambda_S^i\right) \\ & + dt\sum_{S'\neq S}\sum_i B_i\lambda_S^i\rho_{SS'}^i\left[V[K,X,\widetilde{\mathbf{B}}(\mathbf{B}),S'] - V(K,X,\mathbf{B},S)\right] + \left[(1-T_S)KX - I - \gamma I^2\right] dt \end{aligned}$$

The HJB equation states that the risk-neutral expected holding return is equal to the risk-free rate. The second and third lines capture capital gains due to the underlying diffusions in the event of no tax rate change. The final line captures dividends plus capital gains due to tax rate changes. Rearranging terms in

[3]Transitions across Markov matrices drop out, being of order dt^2.

the HJB equation one obtains

$$\begin{aligned}
&\left(r+\sum_i B_i\lambda_S^i\right)V(K,X,\mathbf{B},S) \\
=\; &\max_I \quad V_k(I-\delta K)+V_x(\mu-\sigma\theta)X+\frac{1}{2}\sigma^2X^2V_{xx} \\
&+\sum_j V_{b_j}\left[B_j\left(\sum_i B_i\lambda_S^i-\lambda_S^j\right)+\sum_{i\neq j}B_i\phi_i\pi_{ij}-B_j\phi_j\right] \\
&+\sum_{S'\neq S}\sum_i B_i\lambda_S^i\rho_{SS'}^i V[K,X,\widetilde{\mathbf{B}}(\mathbf{B}),S']+(1-T_S)KX-I-\gamma I^2
\end{aligned} \tag{16}$$

As discussed above, with benefits that are linear in the stock and adjustment costs that are independent of the stock, the value function is separable:

$$V(K,X,\mathbf{B},S)=KQ(X,\mathbf{B},S)+G(X,\mathbf{B},S). \tag{17}$$

Isolating those terms in the HJB equation involving the investment policy I, the optimal instantaneous investment solves:

$$\begin{aligned}
&\max_I \quad Q(X,\mathbf{B},S)I-I-\gamma I^2 \\
\Rightarrow\quad &I_S^*=\frac{Q(X,\mathbf{B},S)-1}{2\gamma};\ S=1,...,N \\
\Rightarrow\quad &I_S^*Q(X,\mathbf{B},S)-I_S^*-\gamma I_S^{*2}=\frac{[Q(X,\mathbf{B},S)-1]^2}{4\gamma}.
\end{aligned} \tag{18}$$

Since the HJB equation must hold pointwise, the terms scaled by K must equate. Using this fact we obtain an equilibrium condition for the shadow value of capital

$$\begin{aligned}
&\left(r+\delta+\sum_i B_i\lambda_S^i\right)Q(X,\mathbf{B},S) \\
=\; &(\mu-\sigma\theta)XQ_x(X,\mathbf{B},S)+\frac{1}{2}\sigma^2X^2Q_{xx}(X,\mathbf{B},S) \\
&+\sum_j\left[B_j\left(\sum_i B_i\lambda_S^i-\lambda_S^j\right)+\sum_{i\neq j}B_i\phi_i\pi_{ij}-B_j\phi_j\right]Q_{b_j}(X,\mathbf{B},S) \\
&+\sum_{S'\neq S}\sum_i B_i\lambda_S^i\rho_{SS'}^iQ(X,\widetilde{\mathbf{B}}(\mathbf{B}),S')+(1-T_S)X.
\end{aligned} \tag{19}$$

The preceding equation states that the expected holding return on capital is equal to the opportunity cost. The holding return consists of dividends plus capital gains associated with the underlying diffusions, along with gains due to tax rate changes.

Since the marginal product of capital is linear in X, we conjecture the shadow value must also be linear in X:

$$Q(X,\mathbf{B},S)=X\Psi_S(\mathbf{B}). \tag{20}$$

Substituting this into the shadow value equation we find that X drops out:

$$\begin{aligned}
&\left(r+\delta-\mu+\sigma\theta+\sum_i B_i\lambda_S^i\right)\Psi_S(\mathbf{B}) \qquad (21)\\
=&\sum_j\left[B_j\left(\sum_i B_i\lambda_S^i-\lambda_S^j\right)+\sum_{i\neq j}B_i\phi_i\pi_{ij}-B_j\phi_j\right]\frac{\partial}{\partial B_j}\Psi_S(\mathbf{B})\\
&+\sum_{S'\neq S}\sum_i B_i\lambda_S^i\rho_{SS'}^i\Psi_{S'}\left(\widetilde{\mathbf{B}}(\mathbf{B})\right)+1-T_S.
\end{aligned}$$

Next, we conjecture that for each of the N states there exists a vector of *shadow value constants* of dimension J solving

$$\Psi_S(\mathbf{B})=\sum_{j=1}^J B_j\Psi_S^j. \qquad (22)$$

That is, each Ψ_S^j allows one to capture the shadow value from the perspective of a hypothetical agent who knows the current instantaneous Markov matrix is j. Under the stated conjecture, pricing is then done taking a belief-weighted average of the j-specific shadow values. Under the maintained conjecture, the shadow value equation (21) can be written as

$$\sum_{j=1}^J B_j\begin{pmatrix}(r+\delta-\mu+\sigma\theta+\lambda_S^j+\phi_j)\Psi_S^j\\ -\lambda_S^j\sum_{S'\neq S}\rho_{SS'}^j\Psi_{S'}^j-(1-T_S)\\ -\phi_j\left(\sum_{i\neq j}\pi_{ji}\Psi_S^i\right)\end{pmatrix}=0. \qquad (23)$$

Since the preceding equation must hold if one sequentially sets each $B_j=1$, we demand that for each $j=1,...,J$ and each state $S=1,...,N$ the bracketed term in the preceding equation must be 0. We then have the following proposition.

Proposition 2. *If tax rate changes are driven by a latent regime shifting Markov chain, the shadow value of capital is*

$$Q(X,\mathbf{B},S)=X\sum_{j=1}^J B_j\Psi_S^j,$$

where the $J\times N$ shadow value constants $\{\Psi_S^j\}$ solve the following system of linear equations

$$\begin{aligned}
1-T_1 &= (r+\delta-\mu+\sigma\theta+\lambda_1^1+\phi_1)\Psi_1^1-\lambda_1^1\sum_{S'\neq 1}\rho_{1S'}^1\Psi_{S'}^1-\phi_1\left(\sum_{i\neq 1}\pi_{1i}\Psi_1^i\right)\\
&\ ...\\
1-T_N &= (r+\delta-\mu+\sigma\theta+\lambda_N^1+\phi_1)\Psi_N^1-\lambda_N^1\sum_{S'\neq N}\rho_{NS'}^1\Psi_{S'}^1-\phi_1\left(\sum_{i\neq 1}\pi_{1i}\Psi_N^i\right)\\
&\ ...\\
1-T_1 &= (r+\delta-\mu+\sigma\theta+\lambda_1^J+\phi_J)\Psi_1^J-\lambda_1^J\sum_{S'\neq 1}\rho_{1S'}^J\Psi_{S'}^J-\phi_J\left(\sum_{i\neq J}\pi_{Ji}\Psi_1^i\right)\\
&\ ...\\
1-T_N &= (r+\delta-\mu+\sigma\theta+\lambda_N^J+\phi_J)\Psi_N^J-\lambda_N^J\sum_{S'\neq N}\rho_{NS'}^J\Psi_{S'}^J-\phi_J\left(\sum_{i\neq J}\pi_{Ji}\Psi_N^i\right).
\end{aligned}$$

It is instructive to compare the determination of shadow values without microeconomic uncertainty (Proposition 1) with the determination of shadow values with microeconomic uncertainty (Proposition 2).

In particular, note that in the special case of Proposition 2 where the underlying Markov matrix is constant over time, with no possibility of regime shifts ($\phi = \mathbf{0}$), the shadow value of capital is determined by taking the shadow values under known constant data generating processes from Proposition 1 and then applying the belief weights to them. That is:

$$\phi = \mathbf{0} \Rightarrow Q(X, \mathbf{B}, S) = \sum_{j=1}^{J} B_j \widetilde{Q}^j(X, S) = X \sum_{j=1}^{J} B_j \widetilde{\Psi}_S^j. \tag{24}$$

With regime shifts, the shadow value constants have a slightly different interpretation. In this case, rather than Ψ_S^j capturing the shadow value when j is known to be the Markov matrix into perpetuity, now Ψ_S^j captures the shadow value from the perspective of a hypothetical agent who knows that at the present instant the stochastic Markov matrix is in regime j.

3.3. Drawing Inferences from Shock Responses

With analytical expressions for shadow values in-hand (Proposition 2), recovering shock responses from causal effects is a simple calculation. To see this, note that the ratio of causal effect to shock response can be written as

$$\frac{CE_{SS'}}{SR_{SS'}} = \frac{\left(\frac{1}{2\gamma}\right)\left(\frac{1}{r+\delta-\mu+\sigma\theta}\right) X_t \times (T_S - T_{S'})}{\left(\frac{1}{2\gamma}\right)\left(Q(X_t, \widetilde{\mathbf{B}}(\mathbf{B}), S') - Q(X_t, \mathbf{B}, S)\right)}. \tag{25}$$

Using Proposition 2 to calculate the denominator in the preceding equation, we obtain a formula for recovering the causal effect implied by a given shock response as shown in the following proposition.

Proposition 3. *The causal effect implied by an observed shock response is*

$$CE_{SS'} = SR_{SS'} \times \frac{(T_S - T_{S'})/(r + \delta - \mu + \sigma\theta)}{\sum_{j=1}^{J} B_j \left[\left(\frac{\lambda_S^j \rho_{SS'}^j}{\sum_i B_i \lambda_S^i \rho_{SS'}^i}\right) \Psi_{S'}^j - \Psi_S^j\right]}. \tag{26}$$

where the shadow value constants $\{\Psi_S^j\}$ are determined per Proposition 2.

A sharper understanding of the determinants of shock responses under model uncertainty is obtained by decomposing them as follows:

$$\begin{aligned}
SR_{SS'} &= \frac{X}{2\gamma}\left[\Psi_{S'}(\widetilde{\mathbf{B}}) - \Psi_S(\mathbf{B})\right] \\
&= \frac{X}{2\gamma}\left[(\Psi_{S'}(\mathbf{B}) - \Psi_S(\mathbf{B})) + \left(\Psi_{S'}(\widetilde{\mathbf{B}}) - \Psi_{S'}(\mathbf{B})\right)\right] \\
&= \frac{X}{2\gamma}\left[\sum_{j=1}^{J}\left(B_j(\Psi_{S'}^j - \Psi_S^j) + \left(\widetilde{B}_j - B_j\right)\Psi_{S'}^j\right)\right] \\
&= \frac{X}{2\gamma}\left[\sum_{j=1}^{J}\left(B_j(\Psi_{S'}^j - \Psi_S^j) + B_j\left(\left(\frac{\lambda_S^j \rho_{SS'}^j}{\sum_i B_i \lambda_S^i \rho_{SS'}^i}\right) - 1\right)\Psi_{S'}^j\right)\right].
\end{aligned} \tag{27}$$

The first term in the preceding equation illustrates that shock responses hinge upon the vector of beliefs held the instant before the tax change arrives. The second term illustrates that shock responses also hinge upon the nature of the belief revision that a specific natural experiment brings about.

It might be hoped that shock response estimates will at least have the same sign as the theory-implied causal effect. However, it is easy to illustrate cases analytically where shock responses have the wrong sign. For example, suppose there is no regime shifting ($\phi = \mathbf{0}$). Suppose also that the current tax state S has the property that for all potential data generating processes, all potential transition-to states (states S' such that $\rho_{SS'}^j > 0$) are absorbing.

With a known Markov matrix and absorbing transition-to states S', we have the following equilibrium condition pinning down shadow values

$$(r+\delta-\mu+\sigma\theta+\lambda_S^j)Q^j(X,S)=\lambda_S^j\sum_{S'\neq S}\rho_{SS'}^j\left(\frac{(1-T_{S'})X}{r+\delta-\mu+\sigma\theta}\right)+(1-T_S)X. \tag{28}$$

From the preceding equation and equation (24) it follows that in the present example

$$Q(X,\mathbf{B},S)=\frac{(1-T_S)X}{(r+\delta-\mu+\sigma\theta)}+\sum_{j=1}^{J}B_j\frac{\lambda_S^j\left[T_S-\sum_{S'\neq S}\rho_{SS'}^jT_{S'}\right]X}{(r+\delta-\mu+\sigma\theta+\lambda_S^j)(r+\delta-\mu+\sigma\theta)}. \tag{29}$$

Thus, with permanent shocks we have

$$\begin{aligned} SR_{S\widetilde{S}} &= \frac{1}{2\gamma}\left[\frac{(1-T_{\widetilde{S}})X}{(r+\delta-\mu+\sigma\theta)}-Q(X,\mathbf{B},S)\right] \qquad (30)\\ &= CE_{S\widetilde{S}}\times\left[1-\sum_{j=1}^{J}B_j\left(\frac{\overbrace{\sum_{S'\neq S}\rho_{SS'}^jT_{S'}-T_S}^{\text{Conditional Expected Change}}}{\underbrace{T_{\widetilde{S}}-T_S}_{\text{Realized Change}}}\right)\left(\frac{\lambda_S^j}{r+\delta-\mu+\sigma\theta+\lambda_S^j}\right)\right]. \end{aligned}$$

The preceding equation implies it is entirely possible that shock responses will not even correctly recover the sign of causal effects. In particular, it is apparent that if agents place sufficiently high probability weights on underlying stochastic processes with a high expected changes (in absolute value), then a relatively small realized change of the same sign will be associated with a shock response opposite in sign to the causal effect. For example, if the waiting time for a corporate tax cut has been long, like President Trump's corporate rate cut, agents might expect a very large tax cut. If only a small rate cut had been delivered, the investment response might well have been negative.

The assumption of permanent shocks is not necessary to generate sign reversals. To see this, consider an economy in which the tax rate has always been high. But suppose that agents think it is possible for tax rates to be cut. In particular, suppose agents know the true latent Markov matrix is fixed ($\phi=\mathbf{0}$) and is one of two types. Markov matrix 1 features a binary tax rate switching between high and medium. Markov matrix 2 features a binary tax rate switching between high and low. For simplicity, assume the shock probability is λdt across all states and across both potential Markov matrices.

Suppose now that the tax rate is cut from high to medium, and consider the shock response. To begin, note that after such a rate change, Bayesian agents will place probability weight 1 on Markov matrix 1. Note also from Proposition 1 it follows that under binary tax rates and a known data generating process (1 or 2), the shadow value constants are

$$\begin{aligned} \begin{bmatrix}\widetilde{\Psi}_H^1\\ \widetilde{\Psi}_M^1\end{bmatrix} &= \begin{bmatrix}\frac{1-T_H}{r+\delta-\mu+\sigma\theta}+\frac{\lambda(T_H-T_M)}{(r+\delta-\mu+\sigma\theta)(r+\delta-\mu+\sigma\theta+2\lambda)}\\ \frac{1-T_M}{r+\delta-\mu+\sigma\theta}+\frac{\lambda(T_M-T_H)}{(r+\delta-\mu+\sigma\theta)(r+\delta-\mu+\sigma\theta+2\lambda)}\end{bmatrix} \qquad (31)\\ \begin{bmatrix}\widetilde{\Psi}_H^2\\ \widetilde{\Psi}_L^2\end{bmatrix} &= \begin{bmatrix}\frac{1-T_H}{r+\delta-\mu+\sigma\theta}+\frac{\lambda(T_H-T_L)}{(r+\delta-\mu+\sigma\theta)(r+\delta-\mu+\sigma\theta+2\lambda)}\\ \frac{1-T_L}{r+\delta-\mu+\sigma\theta}+\frac{\lambda(T_L-T_H)}{(r+\delta-\mu+\sigma\theta)(r+\delta-\mu+\sigma\theta+2\lambda)}\end{bmatrix} \end{aligned}$$

Now let B denote the probability weight placed on Markov matrix 1 prior to the tax rate cut. The shock response here will be

$$SR_{HM} = \frac{1}{2\gamma}\left[Q^1(X,T_M) - (BQ^1(X,T_H) + (1-B)Q^2(X,T_H))\right] \quad (32)$$
$$= \frac{X}{2\gamma}\left[\widetilde{\Psi}^1_M - (B\widetilde{\Psi}^1_H + (1-B)\widetilde{\Psi}^2_H\right]$$
$$= \frac{X}{2\gamma}\left[\frac{(T_H - T_M)(r+\delta-\mu+\sigma\theta+\lambda) - \lambda[T_H - BT_M - (1-B)T_L]}{(r+\delta-\mu+\sigma\theta)(r+\delta-\mu+\sigma\theta+2\lambda)}\right].$$

From the preceding equation it follows

$$\overbrace{1-B}^{\text{Belief Revision}} > \left(\frac{r+\delta-\mu+\sigma\theta}{\lambda}\right)\left(\frac{T_H - T_M}{T_M - T_L}\right) \Rightarrow sgn(SR_{HM}) < 0. \quad (33)$$

That is, the investment response to the tax rate cut will be negative if it brings about a sufficiently negative belief revision. The more general point here is that shock response signs and magnitudes critically depend upon the nature of the belief revision that the tax rate change brings about. In turn, the nature of the belief revision depends upon the specific stochastic environment facing agents.

Hennessy and Strebulaev (2019) analyze natural experiments in dynamic settings with a known shock generating process. They present a simple condition for establishing equivalence between the sign of shock responses and causal effects: *stochastic monotonicity* of the marginal product of capital. If the marginal product of capital is stochastically monotone, then if the marginal product in state S is higher than the marginal product in state S', then at all future dates, the process with initial state S is first-order stochastic dominant to the process with initial state S'. That is, with a known data generating process, stochastic monotonicity ensures that good news today is good news about the future. However, note that in the preceding example, the two potential Markov matrices satisfied stochastic monotonicity respectively, but it was still possible for shock responses to have signs opposite to causal effects. We thus have the following proposition.

Proposition 4. *Stochastic monotonicity of all J potential tax shock generating processes is insufficient to ensure an observed shock response will correctly identify the sign of the theory-implied causal effect.*

Hennessy and Strebulaev (2019) also present a necessary and sufficient condition for shock responses to recover both the sign and magnitude of theory-implied causal effects in a setting with a known data generating process: *martingale marginal product.* Despite the previous proposition's negative result, it turns out that an analogous martingale condition is necessary and sufficient for all potential shock responses to be equal to their respective theory-implied causal effects even in a setting with model uncertainty. To see this, note that if all shock responses are to recover their corresponding causal effect, it must be the case that for all possible states the shadow value of capital must be equivalent to that under permanent tax rates. But from equation (19) if follows that

$$\sum_{S'\neq S}\rho^j_{SS'}T_{S'} = T_S \ \forall\ j \text{ and } \forall\ S \Leftrightarrow Q(X,\mathbf{B},S) = \frac{(1-T_S)X}{r+\delta-\mu+\sigma\theta}\ \forall\ (X,\mathbf{B},S).$$

Thus, we have the following proposition.

Proposition 5. *The necessary and sufficient condition for all potential shock responses to be equal to their respective theory-implied causal effect is that the tax rate be a martingale under all J potential tax shock generating process.*

It is worth stressing that the preceding proposition requires that under *all* potential data generating processes, the tax rate is a martingale. Of course, this will be a demanding condition to satisfy in practice. Nevertheless, this strong condition is necessary to ensure that regardless of current beliefs or the evolution of those beliefs, the tax rate remains a martingale.

Having analyzed the mapping between shock responses and causal effects, we next turn attention to the second potential objective of the econometrician, recovering the investment cost parameter γ from an observed shock response. We know

$$\begin{aligned} SR_{SS'} &= \frac{X}{2\gamma}\left[\Psi_{S'}(\widetilde{\mathbf{B}}) - \Psi_S(\mathbf{B})\right] \\ &\Rightarrow \quad \gamma = \frac{X}{2\times SR_{SS'}}\left[\sum_{j=1}^{J} B_j\left[\left(\frac{\lambda_S^j \rho_{SS'}^j}{\sum_i B_i \lambda_S^i \rho_{SS'}^i}\right)\Psi_{S'}^j - \Psi_S^j\right]\right]. \end{aligned} \tag{34}$$

The preceding equation illustrates that, as was the case with the attempt to recover causal effects from shock responses, correctly recovering deep structural parameters from observed shock responses requires an explicit treatment of the stochastic environment confronting agents–including a specification of the set of possible data generating processes they entertain as possibilities.

A common approach in the public finance literature is to assume agents are completely myopic, in the sense of positing that each tax rate change is viewed as completely unanticipated and permanent. With this approach to imputing shadow values, one would draw an inference $\widehat{\gamma}$ as follows

$$\begin{aligned} SR_{SS'} &= \frac{X}{2\widehat{\gamma}}\left[\frac{1-T_{S'}}{r+\delta-\mu+\sigma\theta} - \frac{1-T_S}{r+\delta-\mu+\sigma\theta}\right] \\ &\Rightarrow \quad \widehat{\gamma} = \frac{X}{2\times SR_{SS'}}\left[\frac{T_S - T_{S'}}{r+\delta-\mu+\sigma\theta}\right] = \gamma \times \frac{CE_{SS'}}{SR_{SS'}}. \end{aligned} \tag{35}$$

The final equality above shows that with the MIT shock assumption, the bias in structural parameter inference is in direct proportion to the bias between shock responses and causal effects.

Consider finally the issue of forecasting the response to a future tax rate change from, say, $T_{S''}$ to $T_{S'''}$ based upon an observed historical shock response to a tax rate change from T_S to $T_{S'}$. Letting B^F and X^F denote the beliefs and aggregate output forecasted at the date of the future tax rate change, it follows from our parameter inference formula (34) that

$$\begin{aligned} SR_{S''S'''} &= \frac{X^F}{2\gamma}\sum_{j=1}^{J} B_j^F\left[\left(\frac{\lambda_{S''}^j \rho_{S''S'''}^j}{\sum_i B_i \lambda_{S''}^i \rho_{S''S'''}^i}\right)\Psi_{S'''}^j - \Psi_{S''}^j\right] \\ &= SR_{SS'} \times \frac{X^F \sum_{j=1}^{J} B_j^F\left(\left(\frac{\lambda_{S''}^j \rho_{S''S'''}^j}{\sum_i B_i \lambda_{S''}^i \rho_{S''S'''}^i}\right)\Psi_{S'''}^j - \Psi_{S''}^j\right)}{X\sum_{j=1}^{J} B_j\left(\left(\frac{\lambda_S^j \rho_{SS'}^j}{\sum_i B_i \lambda_S^i \rho_{SS'}^i}\right)\Psi_{S'}^j - \Psi_S^j\right)}. \end{aligned} \tag{36}$$

Essentially, the preceding formula tells us that correctly extrapolating from a past shock response requires scaling it by the ratio of prospective to historical change in the shadow value of capital. Clearly, as illustrated, extrapolating from past shock responses, even clean shocks, is far from simple. For example, any such forecast is predicated upon making reliable forecasts of future beliefs. But those future beliefs depend upon the precise details of future natural experiments.

4. Numerical Examples

A natural question at this stage is how large is the problem of parameter drift in natural experiments? The objective of this section is to provide calibrated examples based upon historical changes in effective corporate income tax rates.

Consider an econometrician interested in estimating the sign and magnitude of the causal effect of taxes on corporate investment. For the sake of the numerical illustration, assume T_t is the observed history of effective tax rates on corporate investment over the period from 1954-2005, as computed by Gravelle (1994) and the Congressional Research Service (2006).[4]

For the numerical exercises, we discretize the Gravelle/CRS time-series into $S = 3$ tax rate states using the unsupervised machine learning k-means clustering algorithm. Essentially, the k-means algorithm sorts observations into k clusters so as to minimize the Euclidean distance between observed data points and their assigned cluster's centroid. The respective cluster centroids are equal to the within-cluster mean. Applying the k-means algorithm to the Gravelle/CRS tax rate series results in centroid tax rates of 42%, 50% and 58%. With the observed tax rates sorted into their respective clusters, we compute the average transition probability and the average conditional transition probabilities, and then use these as our estimated shock probability and conditional transition probabilities. The resulting time series of tax rate changes between of 42%, 50% and 58% is then used as an input for all of our numerical exercises. The estimated annual tax rate migration matrix is equal to

$$\begin{pmatrix} 0.6929 & 0.3071 & 0.0000 \\ 0.1229 & 0.6929 & 0.1843 \\ 0.0000 & 0.3071 & 0.6929 \end{pmatrix}, \tag{37}$$

where the tax rates are increasing from left to right and from top to bottom.

As shown, we estimate a 30.71% annual probability of a jump in the effective tax rate. This is reflective of the larger number of corporate tax reforms after World War II as well as the fact that changes in inflation led to large changes in effective corporate income tax rates over the sample time period. Two other points are worthy of note in tax rate migration matrix (37). First, there is a slight asymmetry at the 50% tax rate state, with a somewhat higher probability (60%) of a tax rate increase than a tax rate decrease (40%). Second, note that the only positive probability transitions are to nearest neighbor states, and that all transitions are of equal size with $\Delta T = 0.08$.

To complete the model parameterization, we suppose the econometrician inhabits an economy with $r = 2.5\%$ and $\delta = 7.25\%$. These are the same parameter values as used in the numerical examples in Hennessy and Strebulaev (2019). In turn, the real interest rate assumption follows Hennessy and Whited (2005) while the assumed depreciation rate reflects an average of 0 for non-decaying stock variables and the 14.5% depreciation rate assumed by Hennessy and Whited. Alternative γ values would simply change levels of shock responses, whereas our focus below is entirely on relative magnitudes. Finally, following Veronesi (2000) we set the annual instantaneous growth rate of the aggregate output, μ, to 3.3%, the volatility of the aggregate output, σ, to 18%, and the parameter θ to 0.08. Given these parameter values, the theory-implied causal effect for all the shocks considered is $\Delta T/(r + \delta - \mu + \theta\sigma) = 1.0139$. Finally, we limit the number of data generating regimes to two, $J = 2$, and set the switching intensity between them, ϕ, to 0.1 (10 years) in all of our calibration exercises.

We start by considering an economy where nature alternates between two tax rate switching probabilities, $\rho^1_{SS'}$ and $\rho^2_{SS'}$, equal to

$$\rho^1_{SS'} = \begin{pmatrix} 0 & 1 & 0 \\ 0.4 & 0 & 0.6 \\ 0 & 1 & 0 \end{pmatrix} \text{ and } \rho^2_{SS'} = \begin{pmatrix} 0 & 1 & 0 \\ 0.8 & 0 & 0.2 \\ 0 & 1 & 0 \end{pmatrix}, \tag{38}$$

with the tax states ordered as $S = \{42\%, 50\%, 58\%\}$. Note these probability assumptions are consistent with the estimated tax rate migration matrix (37). The tax shock arrival rate λ is set to 0.3071 and is independent of the tax rate state, S, and data generating regime, j.

[Figure 1 about here]

Figure 1 and Table 1 summarize results of this numerical exercise. Both are based upon the assumption that agents enter the economy with initial belief $B_1 = 25\%$. In Figure 1, Panel A shows the evolution of beliefs (blue line), $B_1 = Prob(\rho^j_{SS'} = \rho^1_{SS'})$, and the history of effective tax rates (red line), T_t. Panel B shows Tobin's Q, $Q(X_t, B_1, S)$, scaled by the aggregate output, X_t. Scaling Q by X_t allows us to focus on changes in Q caused solely by changes in tax rates and beliefs. Table 1 quantifies responses of the Q-to-X ratio to changes in tax rates.

[4] This is a simplification because we do not break the total effective tax rate into its constituent parts.

In this simulation exercise changes in the Q-to-X ratio are caused by tax rate changes and by changes in beliefs about the data generating regime, B_1. Agents update their beliefs according to relation (14) only upon observing a tax rate change. In addition, it follows from (38), that only changes from the interim value of 50% to either extreme tax rate value are informative about the data generating process. This is because all probabilities of switching from the extreme tax rate values (42% or 58%) to the interim value of 50% are equal to one under both data generating processes. Indeed, the blue line in Panel A of Figure 1 remains flat in 1962, 1968, 1970, 1976, and 1981, when the tax rate switches to 50%. Since $\rho^1_{21} = 0.4 < \rho^2_{21} = 0.8$, B_1 should discretely jump down upon observing a tax rate reduction from 50% to 42%, and it should jump up upon observing a tax rate hike from 50% to 58%, since $\rho^1_{23} = 0.6 > \rho^2_{23} = 0.2$. Indeed, the blue line in Panel A of Figure 1 jumps down in 1964 and 1982 when the tax rate switches to 42%. Conversely, the blue line jumps up in 1969, 1974, and 1978, when the tax rate switches to 58%. It is also worth mentioning that the Q-to-X ratio jumps discretely since both the tax rates and beliefs jump discretely.

[Table 1 about here]

Table 1 reports changes in the Q-to-X ratio and the corresponding tax rates. The first point worthy of note is that these changes are roughly one-quarter of the theory-implied causal effect equal to 1.0139, a severe downward bias. The second notable point is that while the magnitudes of the responses are different, these differences are relatively small with the maximum difference being 35%. This is mainly due to beliefs not being updated in the absence of tax shocks, a feature of the current data generating process that we alter in our second simulation exercise.

We next consider an economy where nature alternates between two shock arrival intensities $\lambda^1 = 0.0071$ and $\lambda^2 = 0.6071$, both assumed to be independent of the tax rate state, S. This parametrization keeps the average shock arrival intensity equal to 0.3071. The conditional tax rate switching probabilities are given by $\rho^1_{SS'}$ from the first exercise and are set to be the same in both data generating regimes.

Figure 2 and Table 2 summarize results of this numerical exercise. Just like in the previous simulation exercise, both are based upon the assumption that agents enter the economy with initial belief about the data generating regime, $B_1 = Prob(\lambda = \lambda^1)$, equal to 25%. In Figure 2, Panel A shows the evolution of beliefs (blue line), B_1, and the history of effective tax rates (red line), T_t. Panel B shows Tobin's Q, $Q(X_t, B_1, S)$, scaled by the aggregate output, X_t.

[Figure 2 about here]

The first point worthy of note in Figure 2 is that the responses to shocks are all sensitive to waiting time. This is because the beliefs B_1 are evolving over time. Specifically, agents continuously update their beliefs according to (13) in the absence of a tax rate shock. After a tax rate change beliefs exhibit a discrete jump according to (14). For instance, the economy starts in 1954 in the highest tax state with a belief of 25% that the waiting time until a tax reduction will be very long. As time goes by and no tax shock materializes, B_1 sharply increases. Beliefs then experience a large downward jump after the first shock arrives in 1962. Changing beliefs strongly affect the Q-to-X ratio. This is because staying in a highest tax rate state for a long time is "bad news" and, as a result, the Q-to-X ratio falls. Indeed, the Q-to-X ratio declines between 1954 and 1962. By way of contrast, staying in the lowest tax rate state for a long time is "good news" and the Q-to-X ratio should increase if no tax shock occurs. Indeed, Figure 2 shows that in 1982 when the tax rate switches to the lowest tax state, $S = 42\%$, B_1 starts very low and then increases towards its highest value of 85%. The Q-to-X ratio also steadily increases.

The second point worthy of note in Figure 2 is that if the initial tax rate is at one of the extreme values, 42% or 58%, then the magnitude of the response to a shock is very sensitive to waiting time. By way of contrast, if the initial tax rate is at the interim value of 50%, the shock response magnitude is relatively insensitive to waiting time. For instance, the response magnitudes are very similar in 1969 and 1978, while the waiting times are one and two years, respectively. To understand the intuition, notice that, conditional upon a shock arriving, the tax rate change amounts to 8 percentage points if the initial tax rate is at one of the extreme values. By way of contrast, at the intermediate tax rate of 50%, the expected tax rate change, conditional upon a shock arriving, is only 1.6 percentage points. Beliefs about the shock arrival rate are less important if the expected tax rate change, conditional upon a shock, is small.

[Table 2 about here]

Table 2 quantifies responses of the Q-to-X ratio to tax rate changes. Strikingly, Table 2 reveals massive differences in magnitudes of shock responses, despite the fact that all tax rate changes are of equal magnitude and theory-implied causal effects are also of equal magnitude. For example, the minimal shock response has a magnitude of 0.1525 while the maximum shock response magnitude is 0.4241. In other words, the minimum shock response is only 36% of the maximum shock response. This sharply illustrates one of our central points, that historical shock response magnitudes are not generally reliable forecasters of future shock response magnitudes. Nor should they be in economies with learning.

The next point worthy of note in Table 2, related to the first point, is that the magnitude of the response to a first shock has the potential to differ greatly from responses to identical shocks in the future. In this way, the calibrated natural experiment illustrates that causal parameter drift can be quite large in real-world settings. In practice, one could easily envision erroneous dismissals of a first shock response as being a misleading "outlier" inconsistent with "consensus estimates."

Several other points are worth noting in Table 2. First, recall that the theory-implied causal effect for all the shocks considered is 1.0139. However, the magnitude of shock responses never approaches the causal effect. It ranges from about 15% of this value in 1970 to 41% of this value in 1962, a severe downward bias. Second, if agents would have known the data generating process, responses to identical tax rate transitions would be identical. However, with learning it is not the case. For example, the response to a shock in the tax rate from 58% to 50% in 1970 is 0.1525, while the response to an identical tax rate transition in 1981 is 0.2418, a difference of 37%.

5. Macroeconomic Uncertainty

This section extends the baseline model by introducing macroeconomic uncertainty. We follow Veronesi (2000) in assuming the instantaneous drift rate for aggregate output is not observable. One purpose for this extension is to make our framework more realistic and general. However, the primary motivation for this extension is to alert those favoring microeconometric methods to the fact that they must still confront many of the same issues confronting macroeconometricians, even if the tool-kit appears to differ at first glance.

It will be apparent that accounting for macroeconomic uncertainty makes the problem of causal parameter inference in natural experiments even more challenging. Specifically, the correct interpretation of natural experiments hinges upon correctly specifying beliefs about the stochastic processes driving both microeconomic and macroeconomic shocks. Relatedly, while the microeconometric literature seeks to recover unconditional objects, abstracting from macroeconomic state variables, it is apparent that shock responses are functions of both latent and observable macroeconomic state variables.

5.1. Shadow Values Redux

Following Veronesi (2000), the instantaneous drift of aggregate output X can take on any one of $N' \geq 2$ values, $\mu_1 < \mu_2 < ... < \mu_{N'}$. Drifts are indexed by either n or m below. Over any infinitesimal time interval dt with probability pdt a drift will be randomly drawn according to the probability distribution $\mathbf{f} = (f_1, ..., f_{N'})$. Let $\mathbf{Z}$ be the vector of probability weights agents place on each potential drift and let

$$\mu(\mathbf{Z}) \equiv \sum_{n=1}^{N'} Z_n \mu_n. \tag{39}$$

From Lemma 1 in Veronesi (2000) it follows macroeconomic beliefs evolve as a diffusion, with:

$$dZ_n = \underbrace{p(f_n - Z_n)}_{\equiv \mu_{z_n}} dt + \underbrace{\frac{Z_n[\mu_n - \mu(\mathbf{Z})]}{\sigma}}_{\equiv \sigma_{z_n}} dW. \tag{40}$$

Agents are assumed to have identical isoelastic utility functions

$$u(c,t) \equiv e^{-\beta t} \frac{c^{1-\nu}}{1-\nu}. \tag{41}$$

where β is the discount rate and ν is the coefficient of relative risk aversion. The stochastic discount factor (SDF) is

$$M_t \equiv e^{-\beta t} X_t^{-\nu}. \tag{42}$$

As in Cochrane (2001), the risk-free government bond has a constant price of 1 and must therefore pay the following risk-free rate

$$r(\mathbf{Z}) \equiv -\frac{E[dM]}{M} = \beta + \nu\mu(\mathbf{Z}) - \frac{1}{2}\nu(\nu+1)\sigma^2. \tag{43}$$

We now pin down the shadow value of capital, relegating intermediate calculations to the Online Appendix. To begin, the following canonical equilibrium pricing equation must hold for each tax state S:[5]

$$0 = M[(1-T_S)KX - I - \gamma I^2]dt + E_t\{d[MV(K, X, \mathbf{B}, S, \mathbf{Z})]\}. \tag{44}$$

The value function takes the separable form

$$V(K, X, \mathbf{B}, S, \mathbf{Z}) = KQ(X, \mathbf{B}, S, \mathbf{Z}) + G(X, \mathbf{B}, S, \mathbf{Z}). \tag{45}$$

This allows us to rewrite the equilibrium pricing condition as:

$$0 = M[(1-T_S)KX - I - \gamma I^2]dt + E_t\{d(MKQ)\} + E_t\{d(MG)\}. \tag{46}$$

Applying Ito's product rule and dropping terms of order less than dt we have

$$0 = M[(1-T_S)KX - I - \gamma I^2]dt + MQ(I - \delta K)dt + KE_t\{d(MQ)\} + E_t\{d(MG)\}. \tag{47}$$

Isolating those terms in the preceding equation involving the investment control, we find the optimal investment policy takes the standard form

$$\max_I \; M[Q - I - \gamma I^2]dt \Rightarrow I^* = \frac{Q(X, \mathbf{B}, S, \mathbf{Z}) - 1}{2\gamma}. \tag{48}$$

The equilibrium condition must hold on the state space and hence terms scaled by K must equate to zero. Thus, we obtain the following equilibrium condition pinning down the shadow value of capital

$$0 = M(1-T_S)Xdt - \delta MQdt + E_t\{d(MQ)\}. \tag{49}$$

Applying Ito's lemma and dividing by M the previous condition can be restated as:

$$\begin{aligned}
& \left[r(\mathbf{Z}) + \delta + \sum_i B_i \lambda_S^i\right] Q[X, \mathbf{B}, S, \mathbf{Z}] \qquad (50) \\
= \quad & (1-T_S)X + [\mu(\mathbf{Z}) - \nu\sigma^2]XQ_x + \frac{1}{2}\sigma^2 X^2 Q_{xx} \\
& + \sum_j \left[B_j\left(\sum_i B_i \lambda_S^i - \lambda_S^j\right) + \sum_{i \neq j} B_i \phi_i \pi_{ij} - B_j \phi_j\right] Q_{b_j} \\
& + \sum_i B_i \lambda_S^i \sum_{S' \neq S} \rho_{SS'}^i Q[X, \widetilde{\mathbf{B}}(\mathbf{B}), S', \mathbf{Z}] \\
& + \sum_n (\mu_{z_n} - \nu\sigma\sigma_{z_n})Q_{z_n} + \sum_n \sigma\sigma_{z_n} XQ_{xz_n} + \frac{1}{2}\sum_m \sum_n \sigma_{z_m}\sigma_{z_n} Q_{z_m z_n}.
\end{aligned}$$

Notice, this condition is identical to the baseline model's shadow value condition (19) but with the final line added to capture expected capital gains due to the evolution of the macroeconomic belief diffusion processes.

[5]See Cochrane (2001) page 30 for the derivation.

As in the baseline model we conjecture the shadow value is linear in X:

$$Q(X, \mathbf{B}, S, \mathbf{Z}) = X\Psi_S(\mathbf{B}, \mathbf{Z}). \tag{51}$$

Substituting in and simplifying we obtain:

$$\begin{aligned}
&\left[r(Z) + \delta - \mu(\mathbf{Z}) + \nu\sigma^2 + \sum_i B_i\lambda_S^i\right]\Psi_S(\mathbf{B}, \mathbf{Z}) \qquad (52)\\
= \;& (1 - T_S) + \sum_j \left[B_j\left(\sum_i B_i\lambda_S^i - \lambda_S^j\right) + \sum_{i\neq j} B_i\phi_i\pi_{ij} - B_j\phi_j\right]\frac{\partial}{\partial B_j}\Psi_S(\mathbf{B}, \mathbf{Z})\\
&+ \sum_{S'\neq S}\sum_i B_i\lambda_S^i\rho_{SS'}^i\Psi_{S'}[\widetilde{\mathbf{B}}(\mathbf{B}), \mathbf{Z}]\\
&+ \sum_n[\mu_{z_n} + \sigma\sigma_{z_n}(1-\nu)]\frac{\partial}{\partial Z_n}\Psi_S(\mathbf{B}, \mathbf{Z}) + \frac{1}{2}\sum_m\sum_n \sigma_{z_m}\sigma_{z_n}\frac{\partial^2}{\partial Z_m\partial Z_n}\Psi_S(\mathbf{B}, \mathbf{Z})
\end{aligned}$$

Next we conjecture that the shadow value represents a weighted average of microeconomic beliefs as follows:

$$\Psi_S(\mathbf{B}, \mathbf{Z}) = \sum_{j=1}^J B_j\Psi_S^j(\mathbf{Z}). \tag{53}$$

Comparison of equations (22) and (53) is revealing. In the baseline model, each (j, S) shadow value state price Ψ_S^j is a constant. In contrast, with macroeconomic uncertainty, each (j, S) shadow value state price $\Psi_S^j(\mathbf{Z})$ is a function of beliefs about the latent drift.

Substituting the conjectured shadow value function (53) into the shadow value equation (52) and rearranging terms we obtain:

$$\begin{aligned}
&\sum_{j=1}^J B_j\left[\begin{array}{c}\left(r(\mathbf{Z}) + \delta - \mu(\mathbf{Z}) + \nu\sigma^2 + \lambda_S^j + \phi_j\right)\Psi_S^j(\mathbf{Z})\\ -\lambda_S^j\sum_{S'\neq S}\rho_{SS'}^j\Psi_{S'}^j(\mathbf{Z}) - (1 - T_S) - \phi_j\sum_{i\neq j}\pi_{ji}\Psi_S^i(\mathbf{Z})\end{array}\right] \qquad (54)\\
= \;& \sum_{j=1}^J B_j\sum_n[\mu_{z_n} + \sigma\sigma_{z_n}(1-\nu)]\frac{\partial}{\partial Z_n}\Psi_S^j(\mathbf{Z}) + \sum_{j=1}^J B_j\frac{1}{2}\sum_m\sum_n\sigma_{z_m}\sigma_{z_n}\frac{\partial^2}{\partial Z_m\partial Z_n}\Psi_S^j(\mathbf{Z})
\end{aligned}$$

Thus, we demand that for all states S and all potential microeconomic shock generating processes $j = 1, ..., J$:

$$\begin{aligned}
&\left(r(\mathbf{Z}) + \delta - \mu(\mathbf{Z}) + \nu\sigma^2 + \lambda_S^j + \phi_j\right)\Psi_S^j(\mathbf{Z}) \qquad (55)\\
= \;& (1 - T_S) + \lambda_S^j\sum_{S'\neq S}\rho_{SS'}^j\Psi_{S'}^j(\mathbf{Z}) + \phi_j\sum_{i\neq j}\pi_{ji}\Psi_S^i(\mathbf{Z})\\
&+ \sum_n[\mu_{z_n} + \sigma\sigma_{z_n}(1-\nu)]\frac{\partial}{\partial Z_n}\Psi_S^j(\mathbf{Z}) + \frac{1}{2}\sum_m\sum_n\sigma_{z_m}\sigma_{z_n}\frac{\partial^2}{\partial Z_m\partial Z_n}\Psi_S^j(\mathbf{Z}).
\end{aligned}$$

Finally, we conjecture that each (j, S) shadow value state price $\Psi_S^j(\mathbf{Z})$ represents a weighted average over macroeconomic beliefs as follows:

$$\Psi_S^j(\mathbf{Z}) = \sum_{n=1}^N Z_n\Psi_S^{jn}. \tag{56}$$

Essentially, $X\Psi_S^{jn}$ captures shadow value from the perspective of an investor who knows the current instantaneous microeconomic shock process is j and who also knows the current instantaneous drift is μ_n. Under this conjecture we restate our prior condition (55), and now demand that for all states S and all potential microeconomic shock generating processes $j = 1, ..., J$:

$$\sum_{n=1}^{N} Z_n \left[\begin{array}{c} \left[\beta + \delta + \frac{1}{2}\nu(1-\nu)\sigma^2 + p + \lambda_S^j + \phi_j - (1-\nu)\mu_n\right] \Psi_S^{jn} \\ -(1-T_S) - \sum_{S' \neq S} \lambda_S^j \rho_{SS'}^j \Psi_{S'}^{jn} - \left(\sum_{i \neq j} \phi_j \pi_{ji}\right) \Psi_S^{in} \end{array} \right] = p \sum_{m=1}^{N'} f_m \Psi_S^{jm}. \tag{57}$$

Since the right side of the preceding equation does not vary with Z, the term inside brackets must be equal to right side.

We then have the following proposition.

Proposition 6. *If tax rate changes and the drift of aggregate output are driven by latent regime shifting Markov processes then the shadow value of capital is*

$$Q(X, \mathbf{B}, S, \mathbf{Z}) = X \sum_{n=1}^{N'} Z_n \left[\sum_{j=1}^{J} B_j \Psi_S^{jn} \right].$$

where the $J \times N' \times N$ shadow value constants $\{\Psi_S^{jn}\}$ solve the following system of $J \times N' \times N$ linear equations

$$\begin{aligned}
1 - T_1 &= \left[\Gamma - (1-\nu)\mu_1 + \lambda_1^1 + \phi_1\right] \Psi_1^{11} - \lambda_1^1 \sum_{S' \neq 1} \rho_{1S'}^1 \Psi_{S'}^{11} - \phi_1 \sum_{i \neq 1} \pi_{1i} \Psi_1^{i1} - p \sum_{m=1}^{N'} f_m \Psi_1^{1m} \\
&\quad ... \\
1 - T_N &= \left[\Gamma - (1-\nu)\mu_1 + \lambda_N^1 + \phi_1\right] \Psi_N^{11} - \lambda_N^1 \sum_{S' \neq N} \rho_{NS'}^1 \Psi_{S'}^{11} - \phi_1 \sum_{i \neq 1} \pi_{1i} \Psi_N^{i1} - p \sum_{m=1}^{N'} f_m \Psi_N^{1m} \\
&\quad ... \\
1 - T_1 &= \left[\Gamma - (1-\nu)\mu_1 + \lambda_1^J + \phi_J\right] \Psi_1^{J1} - \lambda_1^J \sum_{S' \neq 1} \rho_{1S'}^J \Psi_{S'}^{J1} - \phi_J \sum_{i \neq J} \pi_{Ji} \Psi_1^{i1} - p \sum_{m=1}^{N'} f_m \Psi_1^{Jm} \\
&\quad ... \\
1 - T_N &= \left[\Gamma - (1-\nu)\mu_1 + \lambda_N^J + \phi_J\right] \Psi_N^{J1} - \lambda_N^J \sum_{S' \neq N} \rho_{NS'}^J \Psi_{S'}^{J1} - \phi_J \sum_{i \neq J} \pi_{Ji} \Psi_N^{i1} - p \sum_{m=1}^{N'} f_m \Psi_N^{Jm} \\
&\quad ... \\
1 - T_1 &= \left[\Gamma - (1-\nu)\mu_{N'} + \lambda_1^1 + \phi_1\right] \Psi_1^{1N'} - \lambda_1^1 \sum_{S' \neq 1} \rho_{1S'}^1 \Psi_{S'}^{1N'} - \phi_1 \sum_{i \neq 1} \pi_{1i} \Psi_1^{iN'} - p \sum_{m=1}^{N'} f_m \Psi_1^{1m} \\
&\quad ... \\
1 - T_N &= \left[\Gamma - (1-\nu)\mu_{N'} + \lambda_N^1 + \phi_1\right] \Psi_N^{1N'} - \lambda_N^1 \sum_{S' \neq N} \rho_{NS'}^1 \Psi_{S'}^{1N'} - \phi_1 \sum_{i \neq 1} \pi_{1i} \Psi_N^{iN'} - p \sum_{m=1}^{N'} f_m \Psi_N^{1m} \\
&\quad ... \\
1 - T_1 &= \left[\Gamma - (1-\nu)\mu_{N'} + \lambda_1^J + \phi_J\right] \Psi_1^{JN'} - \lambda_1^J \sum_{S' \neq 1} \rho_{1S'}^J \Psi_{S'}^{JN'} - \phi_J \sum_{i \neq J} \pi_{Ji} \Psi_1^{iN'} - p \sum_{m=1}^{N'} f_m \Psi_1^{Jm} \\
&\quad ... \\
1 - T_N &= \left[\Gamma - (1-\nu)\mu_{N'} + \lambda_N^J + \phi_J\right] \Psi_N^{JN'} - \lambda_N^J \sum_{S' \neq N} \rho_{NS'}^J \Psi_{S'}^{JN'} - \phi_J \sum_{i \neq J} \pi_{Ji} \Psi_N^{iN'} - p \sum_{m=1}^{N'} f_m \Psi_N^{Jm}
\end{aligned}$$

where $\Gamma \equiv \beta + \delta + \nu(1-\nu)\sigma^2 + p$.

Notice, as the linear system is described in the preceding proposition, we first hold fixed the drift at μ_1 and characterize the equilibrium conditions for each microeconomic process j and for each state S. We then let the drift vary up to N'.

As a special case of the preceding proposition, suppose there were no possibility of either microeconomic or macroeconomic regime shifts, with $\phi = \mathbf{0}$ and $p = 0$. In this case, the linear equation system becomes separable into $J \times N'$ distinct blocks of N linear equations, with the solution boiling down to taking a belief weighted average of model solutions under known data generating processes for each combination of microeconomic processes j and drift parameters μ_n. Restated in terms of our tilde notation for known data generating processes, from the preceding proposition and Proposition 1 it follows

$$\phi = \mathbf{0} \text{ and } p = 0 \Rightarrow Q(X, \mathbf{B}, S, \mathbf{Z}) = X \sum_{n=1}^{N'} Z_n \left[\sum_{j=1}^{J} B_j \widetilde{\Psi}_S^{jn} \right]. \tag{58}$$

That is, if there is no regime shifting, one must simply characterize shadow values for each combination of J microeconomic processes and N' potential drifts, as if the model were known, and then apply belief weights, a very simple algorithm. Regime shifting prevents this decomposition, forcing one to invert one relatively large matrix rather than a set of smaller matrices.

5.2. Shock Responses Redux

With the introduction of macroeconomic uncertainty, the ratio of causal effect to shock response is

$$\frac{CE_{SS'}}{SR_{SS'}} = \frac{\left(\frac{1}{2\gamma}\right) X_t \times (T_S - T_{S'}) / [\beta + \delta - (1-\nu)\mu^* + \nu(1-\nu)\sigma^2]}{\left(\frac{1}{2\gamma}\right) \left(Q(X_t, \widetilde{\mathbf{B}}(\mathbf{B}), S', \mathbf{Z}) - Q(X_t, \mathbf{B}, S, \mathbf{Z}) \right)}. \tag{59}$$

Notice, in the preceding equation we are agnostic about the drift the econometrician would like to assume for the purpose of computing the causal effect, and we give it the label μ^*. From the preceding equation it follows that the causal effect implied by an observed shock response is

$$CE_{SS'} = SR_{SS'} \times \frac{(T_S - T_{S'}) / [\beta + \delta - (1-\nu)\mu^* + \nu(1-\nu)\sigma^2]}{\sum_{n=1}^{N'} Z_n \left[\sum_{j=1}^{J} B_j \left(\frac{\lambda_S^j \rho_{SS'}^j}{\sum_i B_i \lambda_S^i \rho_{SS'}^i} \Psi_{S'}^{jn} - \Psi_S^{jn} \right) \right]}. \tag{60}$$

Comparison of the preceding equation with the analogous equation (26) from the baseline model reveals that macroeconomic uncertainty substantially complicates causal inference. Now the econometrician must correctly account for beliefs regarding the aggregate output drift in the denominator. It follows that the magnitude of the wedge between causal effects and shock responses will vary as macroeconomic beliefs vary. Phrased differently, even if one assumed perfect certainty about the underlying process generating the microeconomic shocks, the magnitude of observed responses to identical tax rate shocks would vary considerably with latent macroeconomic beliefs. Given this fact, it is hard to see how any sort of non-contrived consensus could be achieved regarding tax elasticities if that consensus were predicated upon exploiting even ideal exogenous tax rate shocks taking place at different points in time.

The preceding point is best illustrated by way of a numerical simulation. For the purpose of this simulation exercise we consider an economy identical to the one used in the second simulation above but populated by agents with identical isoelastic utility functions. We set the coefficient of relative risk aversion, ν, to be equal to 0.7. In addition to the uncertainty about the tax shock arrival rates, we allow for macroeconomic uncertainty. Specifically, following Veronesi (2000) we assume that over time interval dt with probability $0.5dt$ a drift μ_n is randomly drawn from a pair $\{\mu_1 = 0.075, \mu_2 = 0.005\}$ according to the probability distribution $f = \{0.4, 0.6\}$. The unconditional mean of the drift under the distribution f is equal to 3.3%.

[Figure 3 about here]

Figure 3 and Table 3 summarize results of this numerical exercise. We assume that the initial belief about the microeconomic data generating regime, $B_1 = Prob(\lambda = \lambda^1)$, is equal to 25%. The initial macroeconomic belief is 50%. In Figure 3, Panel A shows the evolution of beliefs (blue line), B_1, and the history of effective tax rates (red line), T_t. Panel B shows Tobin's Q, $Q(X_t, B_1, S)$ scaled by the aggregate output, X_t. It is immediately clear from Figure 3 that macroeconomic uncertainty strongly affects the Q-to-X ratio. For example, the Q-to-X ratio exhibits non-monotone behavior during time intervals between tax rate shocks. However, microeconomic beliefs are strictly monotone during such time intervals. Therefore, the non-monotonicity in the Q-to-X ratio must be driven by time-varying macroeconomic beliefs.

The key point illustrated by this exercise is that uncertainty regarding the macroeconomic data generating process fundamentally alters the magnitude of shock responses. To see this, compare Tables 2 and 3. Every shock response changes. But note, by construction, both tables feature the same microeconomic beliefs at all points in time, since both of them exploit the same time-series of historical tax rates. Therefore, any differences between the respective shock responses across the two tables must be due to the fact that, in Table 3, shock responses are being altered by time-varying macroeconomic beliefs. Phrased differently, the failure to account for macroeconomic uncertainty in Table 3 would lead to faulty inference regarding causal parameters. That is, correctly interpreting the shock responses in Table 3, e.g. mapping them back to theory-implied causal effects would require undoing the confounding effect of both microeconomic and macroeconomic uncertainty, a tall order.

[Table 3 about here]

Comparison of Tables 2 and 3 also reveals that macroeconomic uncertainty can increase the difference between identical shock responses taking place at different points in time. After all, time-varying macroeconomic beliefs can work in the same direction as time-varying microeconomic beliefs to exacerbate shock response differences. For example, in Table 2 which considered a setting without macroeconomic uncertainty, the difference between the 1970 shock response and the identical shock response in 1981 amounted to roughly one-third. However, we see from Table 3, with macroeconomic uncertainty, the difference exceeds 50%. Overall, these simulation results confirm that accounting for macroeconomic uncertainty makes the problem of causal parameter inference in natural experiments even more challenging.

6. 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. In fact, responses to ideally exogenous shocks do not even necessarily clear the low barrier of correct signing of causal effects.

With learning, the correct interpretation of shock responses hinges upon the exact time pattern of realized shocks, as well as (generally unstated) parametric assumptions about priors and potential data generating processes. Conveniently, closed-form formulae were given for: mapping observed shock responses back to theory-implied causal effects; recovering policy-invariant technological parameters; or forecasting future shock responses. Finally, martingale profitability across all potential data generating processes was shown to be a necessary and sufficient condition for shock responses to directly recover comparative statics. However, stochastic monotonicity across all potential data generating processes was shown to be insufficient to ensure shock responses correctly recover the correct sign of theory-implied causal effects.

One final objective of this paper was to formalize concepts and mechanisms that, at present, are either ignored by applied microeconometricians or treated only heuristically. Hopefully, developing a formal framework for the analysis of dynamic natural experiments will clarify points of methodological disagreement between competing camps and facilitate progress through cross-fertilization. Clearly, in many important settings, specifically dynamic settings, the identification challenge mentioned by Heckman (2010) is far from being a settled issue.

References

Abel, A., Eberly, J., 1994. A Unified model of investment under uncertainty. American Economic Review 84, 100-128.

Abel, A., Eberly, J., 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.

Alti, A., 2003. How sensitive is investment to cash flow when financing is frictionless? Journal of Finance 58, 707-722.

Angrist, J.D., Pischke, J.S., 2009. Mostly harmless econometrics: An empiricist's companion. Princeton University Press, Princeton.

Angrist, J.D., Pischke, J.S., 2010. The credibility revolution in economics: How better research design is taking the con out of econometrics. Journal of Economic Perspectives 24, 3-30.

Athey, S., Milgrom, P., Roberts, J., 1998. Robust comparative statics. Working paper, Stanford University.

Bianchi, F., Melosi, L., 2016. Modeling the evolution of expectations and uncertainty in general equilibrium, International Economic Review 57, 717-756.

Bianchi, F., Melosi, L., 2019. Constrained discretion and central bank transparency. Review of Economics and Statistics 100, 187-202.

Bouvard, M., 2014. Real option financing under asymmetric information. Review of Financial Studies 27, 180-210.

Chetty, R., 2012. Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply. Econometrica 80, 969-1018.

Cochrane, J., 2001. Asset pricing. Princeton University Press, Princeton.

Congressional Research Service, 2006. Historical effective tax rates on capital income. CRS Report 1408.

Cummins, J.G., Hassett, K.A., Hubbard, R.G., 1994. A reconsideration of investment behavior using tax reforms as natural experiments. Brookings Papers on Economic Activity 2, 1-74.

Decamps, J., Mariotti, T., 2004. Investment timing and learning externalities. Journal of Economic Theory 118, 80-102.

Goldstein, R., Ju, N., Leland, H., 2001. An EBIT-based model of dynamic capital structure. Journal of Business 74, 483-512.

Gomes, J., 2001. Financing investment. American Economic Review 91, 1263-1285.

Gravelle, J., 1994. The economic effects of taxing capital income. MIT Press, Boston.

Hansen, L., Sargent, T., 2010. Wanting robustness in macroeconomics. Handbook on Monetary Economics 3, 1097-1157.

Heckman, J.J., 2000. Causal parameters and policy analysis in economics: A twentieth century retrospective. Quarterly Journal of Economics 115, 45-97.

Heckman, J.J., Navarro, S., 2007. Dynamic discrete choice and dynamic treatment effects. Journal of Econometrics 136, 341-396.

Hennessy, C.A., Whited, T.M., 2004. Debt dynamics. Journal of Finance 60, 1129-1165.

Hennessy, C.A., Strebulaev, I., 2019. Beyond random assignment: Credible inference and extrapolation in dynamic economies. Journal of Finance, Forthcoming.

Jovanovic, B., 1982. Selection and the evolution of industry. Econometrica 50, 649-670.

Keane, M.P., 2010. Structural vs. atheoretic approaches to econometrics. Journal of Econometrics 156, 3-20.

Keane, M.P., Wolpin, K.I., 2002. Estimating welfare effects consistent with forward-looking behavior. Journal of Human Resources 37, 570-599.

Kocherlakota, N., 2018. Practical policy evaluation. NBER Paper 24643.

Lucas, R.E., Jr., 1976. Econometric policy evaluation: A critique. In: Brunner, K., Meltzer, A. (Eds.), The Phillips curve and labor markets. North-Holland, Amsterdam.

Moyen, N., 2005. Investment-cash flow sensitivities: Constrained vs unconstrained firms. Journal of Finance 59, 2061-2092.

Romer, P., 2016. The trouble with macroeconomics. Available at https://ccl.yale.edu/sites/default/files/files/The%20Trouble%20with%20Macroeconomics.pdf.

Romer, C., Romer, D., 2014. The NBER Monetary Economics Program. NBER Reporter.

Rust, J., 2010. Comments on Structural vs. atheoretic approaches to econometrics. Journal of Econometrics 156, 21-24.

Sims, C., 2010. But economics is not an experimental science. Journal of Economic Perspectives 24, 59-68.

Slemrod, J., 1992. Do taxes matter? Lessons from the 1980's. American Economic Review 82, 250-256.

Summers, L., 1981. Taxation and corporate investment: A q-theory approach. Brookings Papers on Economic Activity 1, 67-132.

Veronesi, P., 2000. How does information quality affect stock returns? Journal of Finance 55, 807-837.

Panel A: Tax rates and beliefs

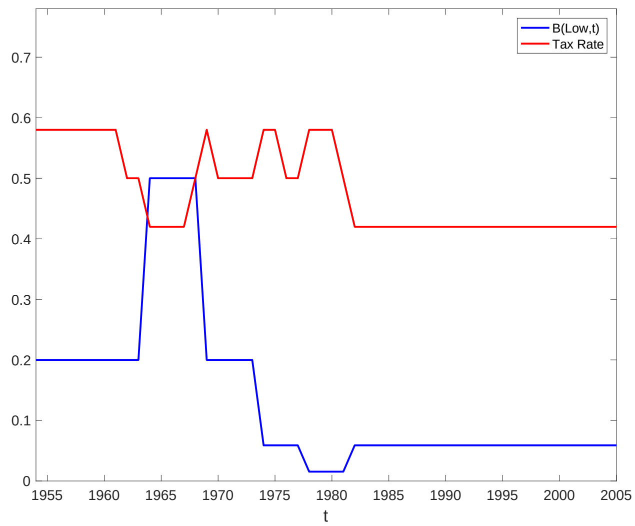

Panel B: Q-to-X ratio

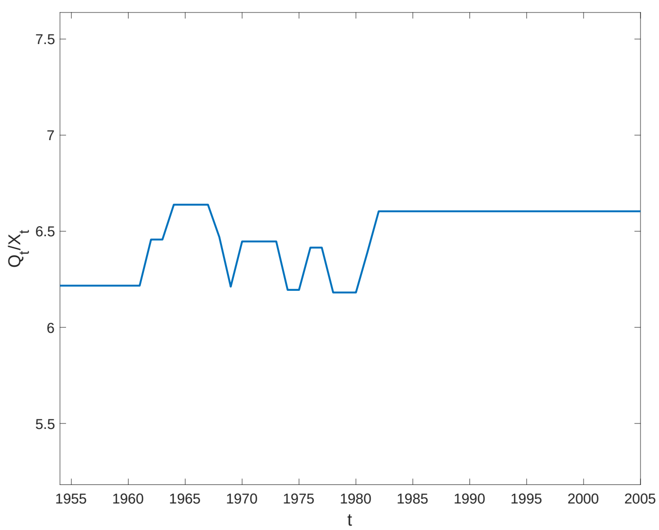

Figure 1 – Simulated Responses to Tax Rate Shocks: Different Switching Probabilities
The figure shows simulated tax shock responses for the case of two different tax rate switching probabilities, $\rho^{1,2}_{SS'}$. Caption of Table 1 provides further details of the simulation. Panel A shows the evolution of beliefs (blue line), $B_1 = Prob(\rho^{j}_{SS'} = \rho^{1}_{SS'})$, and tax rates (red line). Panel B depicts Tobin's Q scaled by the aggregate output, X_t.

Panel A: Tax rates and beliefs

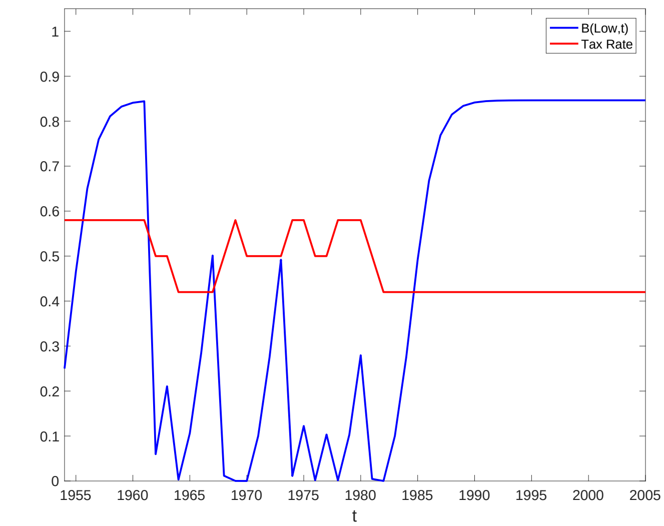

Panel B: Q-to-X ratio

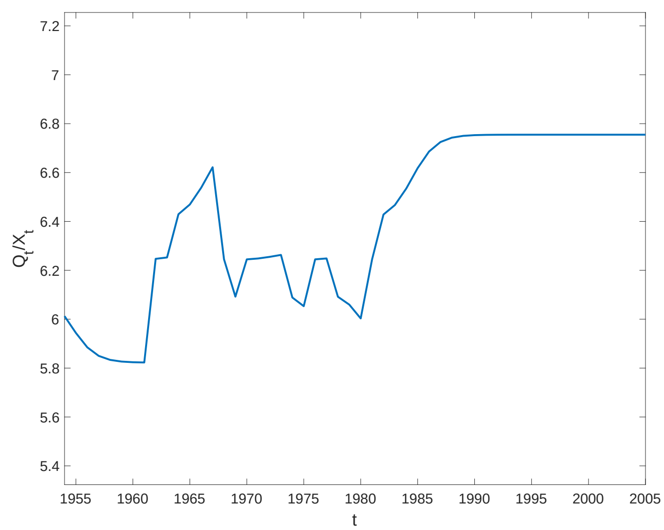

Figure 2 – Simulated Responses to Tax Rate Shocks: Different Shock Arrival Intensities
The figure shows simulated tax shock responses for the case of two different shock arrival intensities, $\lambda^{1,2}$, and the same tax rate switching probabilities, $\rho^1_{SS'} = \rho^2_{SS'}$. Caption of Table 2 provides further details of the simulation. Panel A shows the evolution of beliefs (blue line), $B_1(t) = Prob(\lambda = \lambda^1)$, and tax rates (red line). Panel B depicts Tobin's Q scaled by the aggregate output, X_t.

Panel A: Tax rates and beliefs

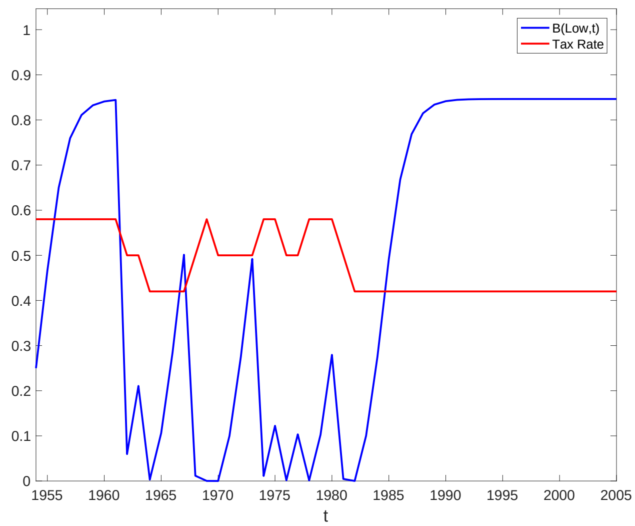

Panel B: Q-to-X ratio

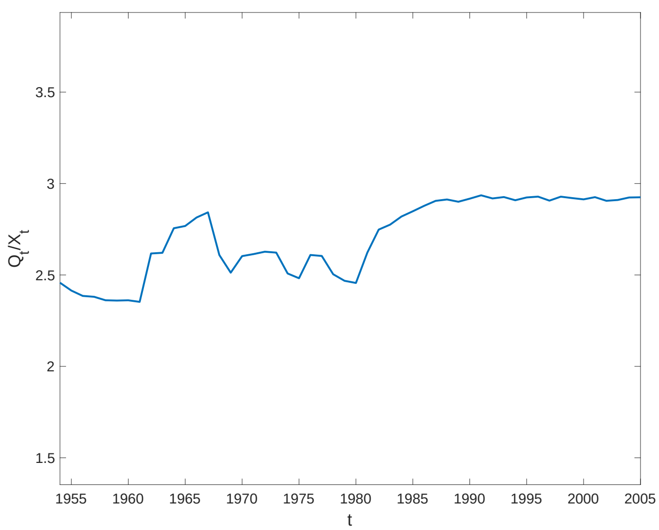

Figure 3 – Simulated Responses to Tax Rate Shocks With Macroeconomic Uncertainty
This figure reports simulated responses to tax rates shock with macroeconomic uncertainty about the instantaneous drift of the aggregate output and microeconomic uncertainty about the tax shock arrival rate. Caption of Table 3 provides further details of the simulation. Panel A shows the evolution of beliefs (blue line), $B_1(t) = Prob(\lambda = \lambda^1)$, and tax rates (red line). Panel B depicts Tobin's Q scaled by the aggregate output, X_t.

Table 1 – Simulated Responses to Tax Rate Shocks: Different Switching Probabilities
This table reports simulated tax shock responses for the case of two different conditional tax rate switching probabilities, $\rho^1_{SS'}$ and $\rho^2_{SS'}$, specified in (38). The historical U.S. 1954-2005 data is used for tax rate shocks with rates alternating between 42%, 50%, and 58%. The tax shock arrival intensity, λ, is set to 0.3071. We report the year of the tax rate shock, change in the Tobin's Q, Q_t, scaled by the aggregate shock, X_t, and the corresponding tax rate.

Year	(1) $\Delta\left(\frac{Q_t}{X_t}\right)$	(2) Tax Rate
1962	0.2399	0.50
1964	0.1814	0.42
1968	-0.1685	0.50
1969	-0.2579	0.58
1970	0.2351	0.50
1974	-0.2519	0.58
1976	0.2199	0.50
1978	-0.2336	0.58
1981	0.2075	0.50
1982	0.2149	0.42

Table 2 – **Simulated Responses to Tax Rate Shocks: Different Shock Arrival Intensities**
This table reports simulated tax shock responses for the case of two shock arrival intensities, $\lambda^1 = 0.0071$ and $\lambda^2 = 0.6071$. The historical U.S. 1954-2005 data is used for tax rate shocks with the tax rate alternating between 42%, 50%, and 58%. The conditional tax rate switching probabilities, $\rho_{SS'}$, with the tax states ordered as $S = \{42\%, 50\%, 58\%\}$, are the same across two data generating regimes and are equal to $\rho^1_{SS'}$ specified in (38). We report the year of the tax rate shock, change in the Tobin's Q, Q_t, scaled by the aggregate shock, X_t, and the corresponding tax rate.

Year	(1) $\Delta\left(\frac{Q_t}{X_t}\right)$	(2) Tax Rate
1962	0.4241	0.50
1964	0.1769	0.42
1968	-0.3765	0.50
1969	-0.1530	0.58
1970	0.1525	0.50
1974	-0.1743	0.58
1976	0.1916	0.50
1978	-0.1568	0.58
1981	0.2418	0.50
1982	0.1833	0.42

Table 3 – **Simulated Responses to Tax Rate Shocks With Macroeconomic Uncertainty**
This table reports simulated responses to tax rates shock with macroeconomic uncertainty about the instantaneous drift of the aggregate output and microeconomic uncertainty about the tax shock arrival rate. The historical U.S. 1954-2005 data is used for tax rate shocks with the tax rate alternating between 42%, 50%, and 58%. The arrival intensities of the tax shocks and conditional transition probabilities for tax rates are the same as reported in the caption of Table 2. Over time interval dt with probability $0.5dt$ a drift μ_n is randomly drawn from a pair $\{\mu_1 = 0.075, \mu_2 = 0.005\}$ according to the probability distribution $f = \{0.4, 0.6\}$. The initial marcoeconomic belief is 50%. The coefficient of relative risk aversion, ν, is set to 0.7. We report the year of the tax rate shock, change in the Tobin's Q, Q_t, scaled by the aggregate shock, X_t, and the corresponding tax rate.

Year	(1) $\Delta\left(\frac{Q_t}{X_t}\right)$	(2) Tax Rate
1962	0.2608	0.50
1964	0.1056	0.42
1968	-0.2372	0.50
1969	-0.0916	0.58
1970	0.0884	0.50
1974	-0.1228	0.58
1976	0.1121	0.50
1978	-0.1123	0.58
1981	0.1826	0.50
1982	0.1132	0.42