Placebo Reforms*

Ran Spiegler†

July 2010

Abstract

An economic variable x evolves according to a stochastic process that is monitored by an infinite sequence of policy makers (PMs). Each PM in his turn decides whether to implement an active reform that may affect the continuation process. Each PM maximizes the change in the value of x that is attributed to him via the following rule: if the PM remains passive, he gets no credit for subsequent changes, whereas if he acts, he gets credit for all subsequent changes until a new reform is implemented. Subgame perfect equilibrium payoffs are uniquely determined by a simple recursive formula, which is subtly related to optimal search models. I use this characterization to illuminate reform timing decisions and risk attitudes of PMs. In particular, when x obeys a linear trend plus *i.i.d* noise, PMs tend to act during crises and display risk aversion in their choice of reform strategies.

1 Introduction

Consider the following hypothetical scenario. You have been appointed as Chief of Police in a certain district. You want to be remembered as someone who brought down crime levels down. As you enter the role, you face a decision whether to implement a large-scale police reform. Although you believe that the reform will lower crime in the long run, you realize that due to short-run fluctuations, things might get worse before they get better. You are concerned that the good effects will be noticeable only after you step down and thus attributed to your successor, while you will take the blame for the short-run downturn.

*Financial support from the Sapir Centre and the ESRC (UK) is gratefully acknowledged.

†Tel Aviv University and University College London. E-mail: r.spiegler@ucl.ac.uk. URL: http://www.homepages.ucl.ac.uk/~uctprsp/

The Chief of Police's predicament is shared by many expert decision makers who care about the perceived outcome of their actions. A surgeon benefits when a patient attributes his recovery to operations the surgeon himself performed. CEOs and sports managers get credit when performance improves shortly after a major recruiting decision. And politicians benefit when GDP growth is perceived as a consequence to their own economic reforms. How should we expect such career/legacy concerns to affect decision makers' actions, particularly when they realize that their successors will face a similar dilemma?

This paper proposes a stylized model of strategic reform choices, in which the objective of each policy maker (PM henceforth) is to get credit for positive changes following his reform decision and to avoid taking the blame for negative changes. In particular, I am interested in how this motive distorts the timing of reforms and the PMs' trade-offs when contemplating various reform strategies.

In the model, an infinite sequence of PMs monitor the stochastic evolution of an economic variable x. Each PM in his turn takes a decision that may affect the continuation of the process that governs x. Although the set of feasible actions is allowed to vary with the state of the process, there is one action, denoted a_0 and interpreted as a "default" or "inaction" option, which is always available to PMs. Each PM maximizes the credit he gets from "the public" for the continuation of the stochastic process, according to the following rule: (i) if the PM selects a_0, he gets no credit for subsequent changes; (ii) when he implements an active reform (which means choosing any $a \neq a_0$) he gets credit for all subsequent changes until a random exogenous "appraisal" takes place, or until a new reform is implemented by another PM - whichever comes first.

Note that in this model, the public's attribution of economic outcomes to PMs' actions is not based on a thorough understanding of the model. Instead, it is based on an intuitive inference rule that traces changes in the economic variable to the most recent intervention. I believe this is a plausible description of the way non-expert outsiders evaluate PMs - e.g., politicians are judged by voters, managers of sports teams are evaluated by a broad fan base, physicians are judged by their patients, etc. In these cases, it makes sense to assume that the evaluator does not have the same level of sophistication as the PMs themselves. This is not an informational asymmetry in the usual sense, but rather an asymmetry in the quality of understanding of the underlying model.

The pitfalls of this type of naive causal linkages when the underlying stochastic processes displays *mean reversion* have been observed by many commentators, notably Kahneman and Tversky (1973). Consider a GP's decision whether to prescribe antibiotics to a patient who displays symptoms that fit both viral and bacterial infection. The timing of the GP's decision is obviously not random - the patient is most likely to turn to the doctor after his health has taken a downturn. If the GP decides to prescribe antibiotics, naive before/after comparison is likely to show recovery and attribute it to the doctor's intervention. This is essentially a *statistical Placebo effect* (as opposed to a truly physiological effect). Indeed, it is often argued (see Goldacre (2009, pp. 38-39) that one of the factors that have led to the growing abuse of antibiotics is the pressure from patients guided by this type of naive inference.

A similar statistical Placebo effect is at play in the context of the present model, where naive public inferences distort PMs' incentives as they contemplate their reform strategies. Specifically, when the process undergoes a transient negative shock, PMs have a stronger incentive to implement an active reform, even when it has no real impact on economic performance. What makes the situation more interesting is the fact that not only the timing of the PM's action, but also the effective timing of his *evaluation*, are biased. Assume that a PM chooses to implement an active reform at time t. Assume further that this reform is later superseded by a new reform implemented by a another PM at time $t' > t$. Then, the former PM can only get credit for changes that took place between t and t'. The fact that the timing t' is determined by the endogenous decisions of future PMs complicates the strategic considerations of the PM who moves at time t.

In Section 3, I characterize subgame perfect equilibria in the infinite-horizon game that the PMs play. Equilibria turn out to have a simple recursive structure: PMs behave as if they collectively solve a dynamic programming problem, in which the value function is

$$V^*(q) = \max_{a \neq a_0} \quad E_{q\prime}[d(q\prime) + \delta \cdot \min(0, V^*(q\prime) \mid (q, a)]$$

where q is the state of the stochastic process that governs the evolution of x; $V^*(q)$ is the maximal payoff that the PM who moves at time t expects to get if he plays $a \neq a_0$ given that the process is in the state q; $d(q)$ is the change in the value of x when the process is in q; and $1 - \delta$ is the probability that an exogenous appraisal occurs in any given period. The function V^* is subtly related to value functions in conventional models of optimal search. Indeed, if all PMs face a choice between two alternatives only - namely, action and inaction - their equilibrium behavior is formally equivalent to a solution of an optimal stopping problem.

This characterization turns out to be very useful when we impose specific classes of stochastic processes. In the main application of the model, analyzed in Section 4, I assume that the stochastic process governing the evolution of x obeys a linear trend with independently distributed noise, where the trend slope and noise distribution are determined by the most recent active reform decision. In this environment, I show that subgame perfect is essentially unique. In each period, PMs choose to implement an active reform if and only if the noise realization drops below a unique, stationary cutoff. Conditional on implementing a reform, the PM chooses an action that maximizes a very simple target function that trades off the expected return from an action and its riskiness. Thus, PMs display risk aversion in equilibrium. When the noise associated with each action is allowed to possess a permanent component, the PMs' equilibrium behavior displays a taste for permanent shocks. These predictions crucially rely on the strategic nature of the model: they disappear in a model with a single PM who acts once and therefore disregards the strategic behavior of his successors.

In Section 5 I turn to a pair of simple examples of other classes of stochastic processes. One example addresses the PMs' trade-off between fundamental reforms with delayed but long-lasting effects and superficial reforms with immediate but short-lived effects. The other example deals with the dynamics of getting out of a crisis. In both cases, I demonstrate that the basic idea of sequentially-moving PMs who are motivated by career/legacy concerns and facing naive public evaluation is capable of illuminating many aspects of real-life strategic reform decisions.

Related literature

My model addresses in an abstract manner a strategic situation that exists in many contexts; it does not commit to a particular application. However, it is closely related to a strand in the political economics literature that deals with the question of reform timing. This literature is primarily preoccupied with explaining why it often seems to be the case that socially beneficial reforms are adopted after a long delay, typically at a time of economic crisis. Drazen and Easterly (2001) provide empirical evidence for this common wisdom. Alesina and Drazen (1991) derive reform delay as a consequence of a war of attrition among interest groups as to which group will bear the burden of reform. Fernandez and Rodrik (1991) explain delay as a form of status quo bias resulting from majority voting when individuals are uncertain about their benefits from reform. In Cukierman and Tommasi (1998), PMs cannot credibly demonstrate the superiority of reform to voters, because the latter are uninformed of the state of the economy and recognize that PMs' policy decisions also reflect their partisan preferences. As a result, socially desirable reforms may fail to be adopted. Orphanides (1992) explains reform delay as a solution to an optimal stopping problem in the context of an inflation stabilization model. At a purely formal level, this work appears to be the closest to the present paper, given the subtle formal relation between my model and stopping problems. For a survey of current approaches to this problem, see Drazen (2001, pp. 403-454).

As a model of a competitive interaction among rational players in the face of agents who employ naive inference rules to evaluate alternatives, this paper is related to Spiegler (2006), a paper that models price competition among providers of credence goods when consumers use anecdotal reasoning to evaluate the quality of each market alternative. The consumers' naive reliance on anecdotal evidence causes them to reward firms for sheer luck as if they had true skill, and as a result a market for an inherently useless product can thrive. This effect is analogous to the statistical Placebo effect that leads the public in the present model to attribute economic performance to PMs' actions.

Finally, the distorting effect of career concerns on experts' intervention decisions was addressed by Fong (2009), who focused on the case of a single PM facing multiple sequential choices, and formulated it as a mechanism design problem of a Bayesian rational evaluator.

2 A Model

An economic variable x evolves over (discrete) time according to the following equation:

$$x^t = x^{t-1} + d^t \tag{1}$$

where the initial condition is $x^0 = 0$, and d^t is governed by the following stochastic process. Let Q be a set of *states*. Let q^t denote the state of the process at time t. For every $q \in Q$, $A(q)$ is a finite set of feasible *actions*. Assume that $|A(q)| \geq 2$ for every $q \in Q$, and that there exists a "null action" a_0, interpreted as *inaction*, such that $a_0 \in \cap_{q \in Q} A(q)$. An action $a \neq a_0$ is interpreted as a reform strategy. In each period $t = 0, 1, 2, ...$, a distinct PM, referred to as player t, observes the entire history of actions $(a_0, ..., a^{t-1})$ and realizations $(x^0, ..., x^t)$, and chooses an action $a^t \in A(q^t)$. Define the *transition function* f as follows: for every $q \in Q$ and $a \in A(q)$, $f(q' \mid q, a)$ is the probability that the process switches to the state $q^{t+1} = q'$, conditional on $(q^t, a^t) = (q, a)$. Finally, define the function $d : Q \to \mathbb{R}$, which records the change in the value of the economic variable when the process is in the state q - i.e., $d^t = d(q^t)$.

To complete this description into a full-fledged infinite-horizon game, we need to describe the players' preferences. Assume that all that PMs care about is the perceived outcome of their actions. Specifically, in each period $t' > t$, player t faces an *appraisal* with exogenous probability $1 - \delta$, where $\delta \in (0, 1)$. An appraisal can occur when the player contends for a new position, or when he reaches an age that prompts an evaluation of his legacy.

Along a path of the game, define $s(t)$ as the *earliest* period $t' > t$ in which either of the following two events occur: *(i)* $a(t') \neq a_0$, or *(ii)* player t faces an appraisal. Then, player t's payoff is

$$\begin{cases} 0 & if \quad a(t) = a_0 \\ x^{s(t)} - x^t & if \quad a(t) \neq a_0 \end{cases}$$

The interpretation of this payoff function is as follows. When a PM remains inactive, none of the subsequent changes in the economic variable are attributed to him. If, on the other hand, the PM implements an active reform, the changes in the economic variable from that moment until another PM implements a new reform - or until his appraisal, whichever comes first - are attributed to him.

Note that the stochastic process itself $(Q, (A(q))_{q \in Q}, f, d)$ is common knowledge among all PMs. Thus, each player t perfectly monitors q^t when he makes his strategy choice. Also, in this section and the next, I assume that Q is finite, to simplify exposition. The extension to infinite state spaces is straightforward.

3 Subgame Perfect Equilibrium

In this section I show that subgame perfect equilibrium payoffs are a well-defined function of the state of the stochastic process. Moreover, this function bears a subtle resemblance to value functions in optimal stopping problems. In subsequent sections I employ this characterization in a number of applications.

We will need to introduce a few pieces of notation at this stage. Define

$$H = \{(q^0)\} \cup \{(q^0, a^0, ..., q^{t-1}, a^{t-1}, q^t)\}_{t=1,2,...}$$

as the set of finite histories in the game. Let $q(h)$ and $t(h)$ denote the state of the process and the identity of the player who moves at the history h. That is:

$$\begin{aligned} q(q^0, a^0, ..., q^{k-1}, a^{k-1}, q^k) &= q^k \\ t(q^0, a^0, ..., q^{k-1}, a^{k-1}, q^k) &= k \end{aligned}$$

For every history h, let (h, a) be the immediate concatenation of h in which player $t(h)$ chooses $a \in A(q(h))$. A strategy profile is denoted $\mathbf{a} = (a(h))_{h \in H}$, where $a(h) \in A(q(h))$. Let $\mathbf{a}(h)$ denote the strategy profile that $\mathbf{a}$ induces in the subgame that begins at the history h.

For a given strategy profile $\mathbf{a}$, define the following function:

$$V_{\mathbf{a}}(h) = \max_{a \in A(q(h)) \setminus \{a_0\}} \quad E[x^{s(t(h))} \mid \mathbf{a}((h, a))] - x^{t(h)}$$

This is simply the maximal expected payoff that player $t(h)$ can attain by choosing to be active. Note that his equilibrium payoff is by definition

$$U_{\mathbf{a}}(h) = \max(0, V_{\mathbf{a}}(h)) \tag{2}$$

because he can guarantee a payoff of zero by choosing the inaction option a_0.

Proposition 1 *Fix a subgame perfect equilibrium* $\mathbf{a}$. *Then,* $q(h) = q(h')$ *implies* $V_{\mathbf{a}}(h) = V_{\mathbf{a}}(h') = V^*(q)$*, where* $V^*(q)$ *is uniquely given by the following recursive equation:*

$$V^*(q) = \max_{a \in A(q) \setminus \{a_0\}} \quad \sum_{q' \in Q} [d(q') + \delta \cdot \min(0, V^*(q'))] \cdot f(q' \mid q, a) \tag{3}$$

Proof. Observe that for every two periods $s > t$:

$$x^s - x^t = (x^s - x^{t+1}) + d^{t+1}$$

Let $\mathbf{a}$ be a subgame perfect equilibrium. Then, at every history h, player $t(h)$ chooses a_0 ($a \neq a_0$) when $V_{\mathbf{a}}(h) < 0$ ($V_{\mathbf{a}}(h) > 0$). By definition, $s(t) = t + 1$ if and only if player $t + 1$ plays $a \neq a_0$. Therefore, we can write $V_{\mathbf{a}}$ recursively as follows:

$$V_{\mathbf{a}}(h) = \max_{a \in A(q(h)) \setminus \{a_0\}} \quad \sum_{(h,a,q')} [d(q') + \delta \cdot \min(0, V_{\mathbf{a}}(h, a, q')] \cdot f(q' \mid q(h), a)$$

The set of histories is countable, and so is the set of values that the R.H.S of this recursive equation can take. The function $V_{\mathbf{a}}$ is a contraction mapping. By the Contraction Mapping Theorem, it has a unique fixed point, such that two histories h and h' with $q(h) = q(h')$ must admit the same value of $V_{\mathbf{a}}$. It follows that $V_{\mathbf{a}}$ is uniquely determined by (3). ■

The recursive equation (3) captures the essence of the PMs' strategic considerations in this model. When player t implements an active reform, he takes into account the future changes in the value of x, but he is concerned that a future PM will act and thus expropriate credit for subsequent changes in the value of x. This future PM will choose to act only if it is profitable to him - i.e., only if the value of V at the time he moves is positive. If this value is negative, the future PM will prefer to be inactive, hence player t will continue to get credit for changes in the value of x.

Comparison with optimal stopping

Recall that for a fixed equilibrium strategy $\mathbf{a}$, the payoff that player $t(h)$ earns is given by (2). Suppose that we replaced the max operator in (3) with a min. Then, the recursive equation that characterizes players' equilibrium payoffs would be

$$U^*(q) = \min\{0 \quad , \quad \min_{a \in A(q) \setminus \{a_0\}} \sum_{q' \in Q} [d(q') + \delta \cdot U^*(q')] \cdot f(q' \mid q, a)\} \tag{4}$$

The function U^* is a value function for a conventional stopping problem (albeit with a non-stationary structure, and where the decision maker has multiple search strategies).

If $A(q) = \{a_0, a_1\}$ for all $q \in Q$, the second min operator in (4) is redundant, and the equilibrium payoffs in our model are indeed reduced to optimal payoffs in a pure stopping problem - i.e., a dynamic decision problem in which the decision maker's only choice is when to stop - such as searching for the lowest price. The difference between our equilibrium behavior in our model and the conventional stopping problem emerges when the action set in some states contains more than two options, and this finds formal expression in the coexistence of max and min operators. The economic significance of this difference will be demonstrated in the applications.

4 Linear Trend Processes with Transient Shocks

In this section, I assume that the set of feasible actions is fixed throughout the game, and denoted A. Every action $a \in A$ is associated with a *trend parameter* μ_a and a continuous density function f_a, which is symmetrically distributed around zero with support $[-k_a, k_a]$. I refer to k_a as the *spread* of f_a. Assume that $\mu_a \in (0, k_a)$ for every $a \in A$ - the role of this assumption is merely to simplify exposition, as it ensures interior solutions - relaxing it would not alter the gist of the analysis. For a given play path, define b^t as the most recent reform strategy that was implemented prior to t. Formally, let t' be the latest period $s < t$ for which $a^s \neq a_0$; then, $b^t = a^{t'}$. If no such period t' exists, set $b^t = a_0$.

The economic variable x evolves according to the following equation:

$$x^t = x^{t-1} + \mu^t + \varepsilon^t - \varepsilon^{t-1} \tag{5}$$

where $\mu^t = \mu_{b^t}$ and ε^t is an independent new draw from f_{b^t}.

Comment: It is straightforward to embed this description in the framework of Section 2. Define the state q^t as a triple $(b^t, \varepsilon^t, \varepsilon^{t-1})$, and let $A(q) = A$ for every q, and $d(q^t) = \mu_{b^t} + \varepsilon^t - \varepsilon^{t-1}$. The transition function is as follows. First, $b^{t+1} = a^t$ if $a^t \neq a_0$, and $b^{t+1} = b^t$ if $a = a_0$. Second, ε^{t+1} is an independent draw from $f_{b^{t+1}}$.

4.1 The Riskiness Function

The following function will play an important role in our analysis. For a given *cdf* F with support $[-k, +k]$, define:

$$R(\varepsilon) = \int_{-k}^{\varepsilon} F(z)dz$$

I refer to R as the *riskiness function* that characterizes the noise distribution associated with the *cdf* F. I use R_a to denote the riskiness function associated with the reform strategy a. In a pair of classic papers, Rothschild and Stiglitz (1970,1971) showed how to use this function to capture the riskiness of a real-valued random variable. In particular, F_b second-order stochastically dominates F_a if and only if $R_a(\varepsilon) \geq R_b(\varepsilon)$ for every $\varepsilon \in (-\infty, +\infty)$.

The following is a useful alternative definition of R:

$$R(\varepsilon) \equiv \varepsilon \cdot F(\varepsilon) - \int_{-k}^{\varepsilon} zf(z)dz \tag{6}$$

Let us recall a few additional properties of the riskiness function.

- $R(0) = 0$ and $R(k) = k$.
- R is strictly increasing with ε.
- $R(\varepsilon) - \varepsilon$ is strictly decreasing with ε.

Finally, the following remark establishes point-wise bounds on the value of the riskiness function.

Remark 1 *For every $\varepsilon \in (-k, +k)$, $\varepsilon < R(\varepsilon) < \frac{1}{2}(\varepsilon + k)$.*

Proof. The L.H.S inequality trivially holds for $\varepsilon \leq 0$. Suppose that $\varepsilon > 0$. We can write

$$\begin{aligned} \varepsilon &= \int_0^{\varepsilon} F(z)dz + \int_0^{\varepsilon} [1 - F(z)]dz \\ R(\varepsilon) &= \int_0^{\varepsilon} F(z)dz + \int_{-\varepsilon}^{0} F(z)dz + \int_{-k}^{-\varepsilon} F(z)dz \end{aligned}$$

By symmetry of the noise, $F(\varepsilon) = 1 - F(-\varepsilon)$ for every $\varepsilon \in [-k, +k]$. Therefore:

$$\int_{-\varepsilon}^{0} F(z)dz = \int_0^{\varepsilon} [1 - F(z)]dz$$

The L.H.S inequality follows immediately.

As to the R.H.S inequality, consider the lottery that assigns probability $\frac{1}{2}$ to each of the two values $\varepsilon = -k$ and $\varepsilon = +k$. Denote the riskiness function induced by the lottery by R^*. This lottery is a mean-preserving spread of every zero-mean noise distribution. Therefore, $R(\varepsilon) \leq R^*(\varepsilon)$ for every $\varepsilon \in [-k, +k]$. Observe that for every $\varepsilon \in (-k, +k)$:

$$R^*(\varepsilon) = \int_{-k}^{\varepsilon} \frac{1}{2} dz = \frac{1}{2}(\varepsilon + k)$$

This establishes the strictness of the R.H.S inequality for interior values of ε. ■

4.2 Equilibrium Analysis

For every $a \in A \backslash \{a_0\}$, define

$$\Pi(a, \varepsilon) = \mu_a - \delta \cdot R_a(\varepsilon)$$

This function trades off the expected trend associated with an active reform strategy and its riskiness. We are now ready for the main result of this section.

Proposition 2 *In any subgame perfect equilibrium, each player t chooses a_0 whenever $\varepsilon^t > \varepsilon^*$, and an action in* $\arg\max_{a \in A \setminus \{a_0\}} \Pi(a, \varepsilon^*)$ *whenever $\varepsilon^t < \varepsilon^*$, where ε^* is uniquely defined by the equation*

$$\max_{a \in A \setminus \{a_0\}} \quad \Pi(a, \varepsilon^*) = (1 - \delta) \cdot \varepsilon^* \tag{7}$$

Proof. Let us write down equation (3), adapted to the present application:

$$V^*(b^t, \varepsilon^{t-1}, \varepsilon^t) = \max_{a \in A \setminus \{a_0\}} \quad \int_{\varepsilon^{t+1}} [\mu_a + \varepsilon^{t+1} - \varepsilon^t + \delta \cdot \min(0, V^*(a, \varepsilon^{t+1}, \varepsilon^t))] \cdot f_a(\varepsilon^{t+1})$$

By assumption,

$$\int \varepsilon f(\varepsilon) = 0$$

hence we can rewrite the above recursive equation as follows:

$$V^*(b^t, \varepsilon^{t-1}, \varepsilon^t) + \varepsilon^t = \max_{a \in A \setminus \{a_0\}} \quad \int_{\varepsilon^{t+1}} [\mu_a + \varepsilon^{t+1} + \delta \cdot \min(0, V^*(a, \varepsilon^{t+1}, \varepsilon^t))] \cdot f_a(\varepsilon^{t+1})$$

Substitute

$$W(b^t, \varepsilon^{t-1}, \varepsilon^t) = V^*(b^t, \varepsilon^{t-1}, \varepsilon^t) + \varepsilon^t$$

so that we can simplify the recursive equation into

$$W(b^t, \varepsilon^{t-1}, \varepsilon^t) = \max_{a \in A \setminus \{a_0\}} \quad \int_{\varepsilon^{t+1}} [\mu_a + \delta \cdot \min(\varepsilon^{t+1}, W(a, \varepsilon^{t+1}, \varepsilon^t))] \cdot f_a(\varepsilon^{t+1})$$

Note that this is an entirely stationary equation, i.e. independent of the state $(b^t, \varepsilon^{t-1}, \varepsilon^t)$. Therefore, its solution is a constant ε^* satisfying:

$$\varepsilon^* = \max_{a \in A \setminus \{a_0\}} \quad \int [\mu_a + \delta \cdot \min(\varepsilon, \varepsilon^*)] \cdot f_a(\varepsilon) \tag{8}$$

Fix a solution $a^* \in A$. Then, we can write this equation as follows:

$$\varepsilon^* = \mu_{a^*} + \varepsilon^* \cdot [1 - F_{a^*}(\varepsilon^*)] + \delta \cdot \int_{-k_{a^*}}^{\varepsilon^*} \varepsilon f_{a^*}(\varepsilon) \tag{9}$$

yielding

$$\Pi(a^*, \varepsilon) = (1 - \delta) \cdot \varepsilon^*$$

The R.H.S (L.H.S) of this equation is strictly increasing (decreasing) in ε^*. By assumption, $\mu_a \in (0, k_a)$ for every a. It follows that there is a unique solution to this equation.

It is straightforward to verify that $a^* \in A \backslash \{a_0\}$ solves the R.H.S of (8) if and only if it maximizes $\Pi(\cdot, \varepsilon^*)$.

Finally, recall that in subgame perfect equilibrium, $a(h) \neq a_0$ if $V_{\mathbf{a}}(h) > 0$ and $a(h) = a_0$ if $V_{\mathbf{a}}(h) < 0$. By the definition of W - and therefore of ε^* - it follows that in equilibrium, each player t chooses $a \in \arg\max_{a \in A}[\mu_a - \delta \cdot R_a(\varepsilon^*)]$ when $\varepsilon^t < \varepsilon^*$ and $a = a_0$ when $\varepsilon^t > \varepsilon^*$. ■

Note that if $\arg\max_{a \in A \backslash \{a_0\}} \Pi(a, \varepsilon)$ is unique for all ε, then the equilibrium is necessarily unique, such that and each player t chooses $a^* = \arg\max_{a \neq a_0} \Pi(a, \varepsilon^*)$ whenever $\varepsilon^t < \varepsilon^*$. In this case, only the first PM who plays $a \neq a_0$ brings a real change in the expected trend, from zero to μ_{a^*}. From that moment, the expected trend is μ_{a^*} forever.

Equilibrium properties

The equilibrium characterization involves two aspects of strategic reform decisions: the timing of reform and the risk attitudes displays in the choice of reform strategies.

Timing. The equilibrium timing of active reform follows a stationary cut-off rule: each player t chooses $a \neq a_0$ if and only if the noise realization in period t does not exceed the cutoff ε^*. Since all noise distributions have a zero mean, the equilibrium expected noise realization conditional on active reform is strictly negative. Thus, the PMs' equilibrium timing of reform gives rise to an "*adverse selection*" effect: the noise realization is negative on average in periods of active reform. In other words, PMs tend to implement reform at times of crisis.

Risk attitudes. Conditional on implementing an active reform, PMs' choices display risk aversion. They choose a reform strategy as if they maximize a utility function that trades off the expected trend and the riskiness of the available reform strategies, in a manner similar to mean-variance preferences, except that the riskiness function R (evaluated at the cutoff ε^*) replaces variance.

The connection to optimal stopping redux

Consider the special case in which the set of reform strategies A is a singleton $\{a_1\}$ - that is, PMs only choose between active reform and inaction. For simplicity, take the $\delta \to 1$ limit. In this case, the model becomes entirely about strategic timing of reform. In equilibrium, each player t chooses $a(t) = a_1$ if and only if $\varepsilon^t \leq \varepsilon^*$, where ε^* is uniquely defined by the equation

$$\mu_1 = R_1(\varepsilon^*)$$

which can be rewritten as

$$\mu_1 = \int_{\varepsilon<\varepsilon^*} (\varepsilon^* - \varepsilon) f_1(\varepsilon)$$

This is precisely the stopping rule in a conventional, stationary optimal stopping problem, in which a consumer, say, searches sequentially for the price of a homogenous product, with a constant per-period search cost. Under this interpretation, μ_1 denotes the search cost, ε denotes the price and ε^* is the optimal cutoff price, below which the consumer always stops the search process and purchases the product. This illustrates the observation made above: in the case of pure reform timing, PMs behave in equilibrium as if they collectively solve a textbook search-for-the-lowest-price problem. This explains why the equilibrium cutoff ε^* increases with μ_1 and decreases when we subject F_1 to a mean-preserving spread.

The effect of a longer horizon

How is equilibrium behavior affected by changes in the PMs' horizon, as captured by the appraisal probability $1 - \delta$?

Corollary 1 *Suppose that the set of reform strategies A is a singleton $\{a_1\}$ - that is, PMs only choose between active reform and inaction. Then, the equilibrium cutoff ε^* decreases with δ.*

To see why this is the case, note that the equilibrium cutoff ε^* satisfies the equation

$$\mu = \delta \cdot [R(\varepsilon^*) - \varepsilon^*] + \varepsilon^* \tag{10}$$

In Section 4.1, we saw that the function $R(\varepsilon) - \varepsilon$ is positive-valued and increasing in ε. Therefore, in order to satisfy (10), ε^* must decrease with δ.

Thus, as PMs' horizon becomes longer (in the sense that their expected appraisal period lies more distantly in the future), they become more reluctant to act, and therefore the equilibrium probability of reform goes down. This result is somewhat surprising, as one might expect that a longer horizon would strengthen the PMs' incentive to act, because the long horizon means that transient noise will be swamped by the overall positive trend. However, recall that the PMs' disincentive to act results from the "adverse selection" that characterizes the timing decisions of future PMs. A longer evaluation horizon raises the probability that the PM will face an appraisal as a result of a subsequent PM's decision to act, rather than from the exogenous process.

Therefore, the disincentive to act becomes stronger. The equivalence to optimal stopping sheds more light on this result: when a consumer searches for a low price, his cutoff price will decrease as he becomes more patient.

When A is not a singleton, the effects of a longer horizon on the equilibrium timing of reform and PMs' risk attitudes become more subtly intertwined, and stronger assumptions are required to obtain clear-cut results. For example, let $A = [l, h] \cup \{a_0\}$, where $h > l > 0$, and assume that for each $a \in [l, h]$, F_a is uniformly distributed over $[-a, +a]$. Thus, each reform strategy is identified with the spread of its noise distribution. In addition, assume that $\mu_a = r \cdot a$ for every $a \in [l, h]$, where $r \in (0, \frac{1}{4})$ is an exogenous constant. In this case, to characterize subgame perfect equilibria, we need to find a noise realization ε^* such that:

$$\max_{a \in [l,h]} \left[ra - \delta \cdot \frac{(\varepsilon^* + a)^2}{4a}\right] = (1 - \delta) \cdot \varepsilon^*$$

The solution to this problem is as follows. For every δ, the equilibrium reform probability is

$$\frac{-(1 - \delta) + \sqrt{1 - \delta(1 - r)}}{\delta} \tag{11}$$

It can be verified that this expression increases with r and decreases with δ.

Let us turn to the PMs' choice of reform strategy conditional on acting. When $\delta < 4r$, all PMs choose the action h in equilibrium after every history. When $\delta > 4r$, all PMs choose the action l after every history. Thus, when the appraisal probability is low (high), PMs opt for the lowest-risk, lowest-return (highest-risk, highest-return) reform strategy.

Ex-ante payoffs

Suppose that just before player t observes ε^t, he is asked to evaluate his equilibrium expected payoff. Since ε^t is the only payoff-relevant aspect of a state $(b^t, \varepsilon^t, \varepsilon^{t-1})$, we can write V^* as a function of ε^t, without loss of generality. For simplicity, take the $\delta \to 1$ limit, and assume that $\arg\max_{a \in A \setminus \{a_0\}} \Pi(a, \varepsilon^*) = \{a^*\}$. Then, the player's ex-ante expected payoff is

$$\int \max(0, V(\varepsilon)) f_{a^*}(\varepsilon)$$

But we know that $V(\varepsilon^*) = 0$. Therefore, $V(\varepsilon) = 0$ for all $\varepsilon > \varepsilon^*$ and $V(\varepsilon) = \varepsilon^* - \varepsilon$ for all $\varepsilon < \varepsilon^*$, hence the ex-ante expected payoff is

$$\int_{-\infty}^{\varepsilon^*} (\varepsilon^* - \varepsilon) f_{a^*}(\varepsilon) = R_{a^*}(\varepsilon) = \mu_{a^*}$$

That is, the ex-ante expected payoff is equal to the trend parameter that characterizes the active reform strategy taken in equilibrium.

Comparison with a Non-Strategic Model

To get a deeper understanding of the strategic considerations of PMs in our model, it will be useful to draw a comparison with a simpler model, in which there is a single PM who acts at an arbitrary period t, and does not expect subsequent PMs to act. The PM's payoff function is exactly as in the model presented in period 2. However, the period $s(t)$ is determined only by the exogenous appraisal process.

The PM's expected payoff from taking an action $a \neq a_0$ in period t, given x^t and ε^t, is

$$\frac{\mu_a}{1-\delta} - \varepsilon^t \tag{12}$$

The reason is simple. The expected number of periods that elapse until the PM faces an appraisal is $\frac{1}{1-\delta}$. Moreover, since the appraisal process is entirely exogenous, it lacks the "adverse selection" effect we encountered in the strategic model, such that the expected noise in the appraisal period is zero. Thus, player t will act if and only if $\varepsilon^t \leq \tilde{\varepsilon}$, where $\tilde{\varepsilon}$ is given by

$$\max_{a \in A} \mu_a = (1-\delta)\tilde{\varepsilon} \tag{13}$$

Conditional on acting, he will choose a to maximize μ_a.

Compare this with the PMs' equilibrium behavior in our model, given by (7). The two equations are nearly identical, except for the term $\delta \cdot R_a(\varepsilon)$, which appears in the strategic model only. This term is crucial, as it leads to several notable differences between the cutoff rules in the two models.

- The cutoff (and therefore the reform probability) is lower in the strategic model. In particular, when $\max_{a \in A} \mu_a \approx 0$, the cutoff is $\varepsilon^* \approx 0$ in the non-strategic model, whereas it is typically negative and bounded away from zero in the strategic model. The reason is that only in the strategic model, PMs are concerned with the "adverse selection" that characterizes the noise realization when a future PM implements an active reform.

- The noise variable is irrelevant for the PM's decision in the non-strategic model, whereas it plays a crucial role in the PMs' equilibrium decisions in the strategic model. The reason, once again, is that the "adverse selection" effect exists only in the strategic model. In the absence of this effect, PMs do not care about the noise because they are risk-neutral. In the strategic model, they care about the noise because it affects the magnitude of the "adverse selection" effect.

- The effect of extending the PMs' horizon is different in the two models. In the non-strategic model, a higher δ leads to a higher reform probability (because we assumed $\mu_a > 0$ for all a). In the $\delta \to 1$ limit, PMs almost always act. In contrast, in the strategic model, as we saw, a higher δ can result in a lower equilibrium reform probability because it makes the strategic "adverse selection" effect more important for PMs.

4.3 Permanent Shocks

In this sub-section, I extend the linear trend model of this section by incorporating the possibility of permanent shocks. Specifically, the process (5) is modified as follows:

$$x^t = \begin{cases} x^{t-1} + \mu^t + \varepsilon^t & \text{with probability } \rho \\ x^{t-1} + \mu^t + \varepsilon^t - \varepsilon^{t-1} & \text{with probability } 1-\rho \end{cases}$$

Thus, ρ is the probability that the random, independent shock in any given period is permanent.

The characterization of subgame perfect equilibrium undergoes a slight modification under this extension. For every $a \in A \backslash \{a_0\}$, define

$$\mu_a^* = \frac{\mu_a}{1-\rho}$$

As before, to ensure an interior solution, assume $\mu_a^* \in (0, k_a)$ for every $a \in A \backslash \{a_0\}$.

Proposition 3 *Subgame perfect equilibrium is characterized exactly as in Proposition 2, except that μ_a^* substitutes μ_a.*

I omit the proof, as it follows the same outline as the proof of Proposition 2. It immediately follows from this characterization that in equilibrium, all PMs display a preference for reform strategies that are associated with permanent shocks. To put it figuratively, suppose that reform involves appointing a new administrator. The impact of any new administrator on the economic variable is uncertain. However, it is likely to be durable if the administrator is appointed for a long term. Then, in equilibrium, PMs prefer reforms that involve such long-term appointments. The intuition for this result is simple: when shocks are more likely to be permanent, the mean reversion that causes the adverse selection effect on subsequent PMs' timing decisions is attenuated, and therefore the impediment to active reform shrinks. It can be shown that when ρ is sufficiently large, all PMs choose an action in $\arg\max_{a \neq a_0} \mu_a$ after every history - i.e., regardless of the noise realization.

4.4 Non-Stationary Processes

The linear trend model analyzed so far is stationary, as the continuation game faced by each PM is identical. In this sub-section I continue to assume that the economic variable evolves according to a linear trend with independent noise. At the same time, I relax the stationarity assumption and allow the set of feasible trend parameters to be history-dependent. Throughout this subsection, I focus on the $\delta \to 1$ limit.

Formally, let S be a finite set of *trend states*. The state of the process in period t is a triple $q^t = (s^t, \varepsilon^t, \varepsilon^{t-1})$. Every trend state s is associated with a trend parameter μ_s, such that $d(q^t) = \mu_s + \varepsilon^t - \varepsilon^{t-1}$. The set of feasible actions at a state is entirely determined by the trend state. In what follows, we will use $A(s)$ to denote the action set whenever the system is at the trend state s. Assume that $a_0 \in A(s)$ and $|A(s)| \geq 2$ for every $s \in S$. The transition function consists of two components. First, transition between trend states is governed by a deterministic transition function τ^*, such that for every $s \in S$ and every $a \in A(s)$, $\tau^*(s,a) \in S$. In particular, $\tau^*(s, a_0) = s$ for every $s \in S$. Second, the noise realization ε^t is *i.i.d* according to a density function f, which is symmetrically distributed around zero with support $[-k, k]$.

We will say that the trend state $r \neq s$ is *immediately reachable* from the trend state s if there exists an action $a \in A(s)$ such that $\tau^*(s,a) = r$. Let $s(t)$ denote the trend state at period t along some given play path.

Proposition 4 *There is a unique collection of cutoffs $(\varepsilon_s^*)_{s \in S}$ such that in any subgame perfect equilibrium, every player t PM chooses a_0 whenever $\varepsilon^t > \varepsilon^*_{s(t)}$, and an action $a \in A(s) \backslash \{a_0\}$ that maximizes*

$$\mu_{\tau^*(s,a)} + \varepsilon^*_{\tau^*(s,a)} - R(\varepsilon^*_{\tau^*(s,a)}) \tag{14}$$

whenever $\varepsilon^t < \varepsilon^*_{s(t)}$. Moreover,the collection $(\varepsilon_s^*)_{s \in S}$ satisfies

$$\varepsilon^*_s \geq \mu_r + \varepsilon^*_r - R(\varepsilon^*_r)$$

for every trend state r that is immediately reachable from s. This inequality is binding when $r = \tau(s, a^*(s))$, where $a^*(s)$ solves (14).

Proof. Let us write down equation (3), adapted to the present application:

$$V^*(s, \varepsilon^{t-1}, \varepsilon^t) = \max_{a \in A(s) \setminus \{a_0\}} \int_{\varepsilon^{t+1}} [\mu_{\tau^*(s,a)} + \varepsilon^{t+1} - \varepsilon^t + \delta \cdot \min(0, V^*(\tau^*(s,a), \varepsilon^{t+1}, \varepsilon^t))] \cdot f(\varepsilon^{t+1})$$

By assumption,

$$\int \varepsilon f(\varepsilon) = 0$$

hence we can rewrite the above recursive equation as follows:

$$V^*(s, \varepsilon^{t-1}, \varepsilon^t) + \varepsilon^t = \max_{a \in A(s) \setminus \{a_0\}} \int_{\varepsilon^{t+1}} [\mu_{\tau^*(s,a)} + \varepsilon^{t+1} + \delta \cdot \min(0, V^*(\tau^*(s,a), \varepsilon^{t+1}, \varepsilon^t))] \cdot f(\varepsilon^{t+1})$$

Substitute

$$W(s, \varepsilon^{t-1}, \varepsilon^t) = V^*(s, \varepsilon^{t-1}, \varepsilon^t) + \varepsilon^t$$

so that we can simplify the recursive equation into

$$W(s, \varepsilon^{t-1}, \varepsilon^t) = \max_{a \in A(s) \setminus \{a_0\}} \int_{\varepsilon^{t+1}} [\mu_{\tau^*(s,a)} + \delta \cdot \min(\varepsilon^{t+1}, W(\tau^*(s,a), \varepsilon^{t+1}, \varepsilon^t))] \cdot f(\varepsilon^{t+1})$$

Note that the solution to this equation is independent of ε^{t-1} and ε^t. Hence, we can write it as a function of the trend state. Thus, we have a system of equations:

$$\varepsilon^*(s) = \max_{a \in A(s) \setminus \{a_0\}} \int [\mu_{\tau^*(s,a)} + \delta \cdot \min(\varepsilon, \varepsilon^*(s))] \cdot f(\varepsilon) \tag{15}$$

Fix a solution $(a^*(s))_{s \in S}$. Then, we can write this equation as follows:

$$\varepsilon^*_s = \mu_{\tau^*(s,a)} + \delta \cdot \int_{-k}^{\varepsilon^*(\tau^*(s,a))} \varepsilon f(\varepsilon) + \varepsilon^*(\tau^*(s,a)) \cdot [1 - F(\varepsilon^*(\tau^*(s,a)))] \tag{16}$$

Set $\delta = 1$ and $r = \tau^*(s,a)$. Then:

$$\varepsilon^*_s = \mu_r + \varepsilon^*_r - R(\varepsilon^*_r) \tag{17}$$

Finally, just as in the stationary model, ε^*_s is a cutoff: whenever $s^t = s$, player t chooses $a \in \arg\max_{a \in A(s) \setminus a_0} [\mu_{\tau^*(s,a)} + \varepsilon^*_{\tau^*(s,a)} - R(\varepsilon^*_{\tau^*(s,a)})]$ when $\varepsilon^t < \varepsilon^*_s$ and $a = a_0$ when $\varepsilon^t > \varepsilon^*_s$. ■

This result has a simple corollary for the relation between average changes in the economic variable and the riskiness function. Note that since the transition between trend states is deterministic, we can pin down the trend state s^t in every period t along a subgame perfect equilibrium play path.

Corollary 2 *The following equation holds i every subgame perfect equilibrium path:*

$$\lim_{T\to\infty}\frac{1}{T}\sum_{t=1}^{T}\mu_{s^t}=\lim_{T\to\infty}\frac{1}{T}\sum_{t=1}^{T}R(\varepsilon^*_{s^t}) \tag{18}$$

This follows immediately from the system of equations (17). Along an equilibrium path, $\varepsilon^*_{s^t}=\mu_{s^{t+1}}+\varepsilon^*_{s^{t+1}}-R(\varepsilon^*_{s^{t+1}})$. Summing over t, we obtain (18). Thus, the relation between the trend parameter and the riskiness function, derived in the stationary model, extends to the non-stationary model, except that the relation is true only in a long-run average sense.

5 Two Examples

The formalism introduced in Sections 2 and 3 is defined for general underlying processes that govern the economic variable. The main application, analyzed in Section 4, was to linear trend processes. In this section I use two simple examples to demonstrate the formalism's wider scope of applications.

5.1 Delayed Effects of Reform

The driving force in the model is the PMs' concern that a future PM will get credit for improvements in the economic variable brought about by their own actions. Therefore, we should expect PMs to prefer superficial reforms that generate immediate benefits, even if these are substantially lower than the benefits from fundamental reforms. In this sub-section I construct a simple application of the model that captures this intuition.

I use the formalism of Section 4.4. The set of trend states is $S=\{s_0,s_1,s_2,s_3,s_4\}$. Let $A(s)=\{a_0,a_L,a_H\}$ for every s, where a_1 represents a superficial reform and a_2 represents a fundamental reform. The trend parameters $(\mu_s)_{s\in S}$ and the transition function τ^* are given by the following table:

s	μ_s	$\tau^*(s,\cdot)$
s_0	0	$a_0 \rightarrow s_0$ $a_L \rightarrow s_1$ $a_H \rightarrow s_2$
s_1	μ_L	$a_0, a_1, a_2 \rightarrow s_3$
s_2	0	$a_0, a_1, a_2 \rightarrow s_4$
s_3	0	$a_0, a_1, a_2 \rightarrow s_3$
s_4	μ_H	$a_0, a_1, a_2 \rightarrow s_4$

where $0 < \mu_L, \mu_H < k$. In period 0, the process is in the trend state s_0. The transient noise ε^t is $i.i.d$ according to the density function f, which is symmetrically and continuously distributed around zero with support $[-k.k]$.

The deterministic process that governs the evolution of the trend state has a simple interpretation. Initially, the trend parameter is zero, and remains so as long as PMs are inactive. If a PM chooses a superficial reform, the trend parameter jumps to μ_L in the next period, but then switches to an absorbing zero-trend state. In contrast, if a PM chooses a fundamental reform at the initial trend state, the process first switches to another zero-trend state but then jumps to an absorbing trend state with μ_H.

Proposition 5 *In subgame perfect equilibrium, the trend states s_2 and s_4 are never reached, and the process reaches the absorbing trend state s_3 in finite time.*

Proof. Suppose that the process is in the trend state s_0 in period t, and that player t chooses a_2. The process then switches to the trend state s_2 in the next period. If player $t+1$ then chooses $a \neq a_0$, player t's payoff is ε. On the other hand, if player $t+1$ chooses a_0 - an event that occurs with probability $1 - F(\varepsilon^*)$ - player t's continuation payoff is

$$\frac{\mu_H + \int_{-k}^{\varepsilon^*} \varepsilon f(\varepsilon)}{F(\varepsilon^*)}$$

by our characterization of equilibrium in the linear trend model of Section 4. Thus, player t's expected payoff is

$$\begin{aligned} & \int_{-k}^{\varepsilon^*} \varepsilon f(\varepsilon) + [1 - F(\varepsilon^*)] \cdot \frac{\mu_H + \int_{-k}^{\varepsilon^*} \varepsilon f(\varepsilon)}{F(\varepsilon^*)} - \varepsilon^t \\ = \quad & \varepsilon^* - R(\varepsilon^*) - \varepsilon^t < -\varepsilon^t \end{aligned}$$

Now suppose that when the process is in the trend state s_0 in period t, player t chooses a_1. Then, his continuation payoff is

$$\mu_L - \varepsilon^t > -\varepsilon^t$$

The reason is that once the process reaches s_1, all players choose a_0, by our characterization of equilibrium in the linear trend model of Section 4. It follows that if player t chooses to be active, he necessarily selects a_1 over a_2. Since $\varepsilon^t < \mu_L$ with positive probability, the process will switch from s_0 to s_1 in finite time, and thus reach the absorbing state s_3 in finite time. ■

Thus, PMs sometimes implement a superficial reform in equilibrium, as long as the noise realization is sufficiently low, but they never implement the fundamental reform because of its delayed effect.

The result, although not surprising, is not entirely trivial. If we assumed that each player plays a_0 with some fixed exogenous probability λ, then a PM's choice between fundamental and superficial reform would depend on the parameters μ_L, μ_H. However, since λ is determined endogenously, we obtain the extreme result that in equilibrium, PMs always prefer superficial reform, regardless of the values of μ_L and μ_H.

5.2 Getting in and out of a Crisis

In this sub-section I explore a simple example that illuminates the dynamics of getting in and out of an economic crisis when PMs are motivated by career/legacy concerns. Suppose that the economy can be in three possible "level" states: normal, mild crisis and deep crisis. As the game begins, the economy is in a state of mild crisis. If an active reform is immediately implemented, the system returns back to normal and remains in this state indefinitely. If an active reform is not implemented, the system dips further to a state of deep crisis, and stays there until an active reform is implemented, in which case the system returns to the mild crisis state, and so forth.

Note that the notion of a state in this description is different from the one adopted elsewhere in the paper, because states are identified with *levels* of the economic variable, rather than with changes. This calls for a slight modification of our model. The state of the system in period t is a triple $(q^*, \varepsilon^t, \varepsilon^{t-1})$, where q^* is a level state and ε^t is a noise realization in period t. For each level state q^*, let $l(q^*) \in \{0, -d, -2d\}$. If the system is in level state q^* in period t, then $x^t = l(q^*) + \varepsilon^t - \varepsilon^{t-1}$. The noise realization ε^t is *i.i.d* according to a density function f, which is symmetrically and continuously distributed around 0 with support $[-k, k]$. The transitions between level states are deterministic, and given by the following table:

q^*	$l(q^*)$	$t(q^*, \cdot)$
q_0^*	$-d$	$a_1 \to q_2^*$
		$a_0 \to q_1^*$
q_1^*	$-2d$	$a_1 \to q_0^*$
		$a_0 \to q_1^*$
q_2^*	0	$a_0, a_1 \to q_2^*$

It can be shown that subgame perfect equilibrium has exactly the same characterization as in Proposition 4, except that for every pair of level states s, r such that r is immediately reachable from s, μ_r should be replaced by $l(r) - l(s)$. This leads to the following result.

Proposition 6 *In subgame perfect equilibrium: (i) when the system is in level state q_0^* in period t, player t chooses a_1 when $\varepsilon^t > d$ and a_0 when $\varepsilon^t < d$; (ii) when the system is in level state q_1^* in period t, player t chooses a_1 when $\varepsilon^t > \varepsilon_2^*$ and a_0 when $\varepsilon^t < \varepsilon_2^*$, where $\varepsilon_2^* = 2d - R(d)$; (iii) when the system is in the level state q_2^*, PMs always choose a_0.*

Proof. When the system reaches q_2^*, no PM ever chooses a_1 because from that moment, $E(x^{t+1} - x^t) = 0$ for every period t - i.e., it as if we are in the stationary linear trend model of Section 4 with a zero trend coefficient. Therefore, when the system reaches q_1^*, PMs will choose a_1 whenever $\varepsilon^t < \varepsilon_1^*$, and $\varepsilon_1^* = d$. Finally, when the system reaches q_2^*, PMs will choose a_1 whenever $\varepsilon^t > \varepsilon_2^*$, where ε_2^* is given by equation (17), adapted to the present example:

$$\varepsilon_2^* = d + \varepsilon_1^* - R(\varepsilon_1^*) = 2d - R(d)$$

■

Note that $\varepsilon_2^* < d$. This result captures a spiral of plunging into an ever deeper crisis. As the crisis worsens, PMs become more reluctant to act, and a more negative transient shock is required to prompt them to act. The intuition is as follows. When the system is initially in a mild crisis, PMs do not face the adverse selection consideration because if they act, the system switches to the normal state indefinitely and from that moment no PM ever acts. Therefore, as long as the noise realization is not too positive (i.e., not above d), PMs will act. However, if the noise realization during a mild crisis is sufficiently positive (above d), it masks the structural decline and therefore PMs refuse to act. But this causes further deterioration into deep crisis. The problem is that in that state, PMs *do* face the adverse selection consideration and are therefore more strongly deterred from acting.

6 Conclusion

My objective in this paper was to present a simple model of strategic policy making when PMs have career/legacy concerns and they are evaluated according to a naive causal-inference rule. The merit of this model is that it illuminates the much-researched subject of reform delay from a new angle, which links reform delay to the risk and intertemporal preferences exhibited by PMs' equilibrium reform decisions.

The model's very simplicity immediately suggests various extensions that were not examined here for the sake of brevity. First, throughout the paper, we assumed that active reform strategies are costless for PMs. It would be interesting to see how imposing a cost structure on the set of reform strategies would affect equilibrium analysis. Second, I have modelled the PMs' horizon in a very simplistic way, capturing it with the single parameter δ. There are natural ways to complicate this component of the model. For example, the value of δ for each PM could be stochastic. Alternatively, the possibility of multiple terms and re-election could be modeled explicitly. Third, I have assumed that PMs have nothing but career concerns, and one may wish to explore what happens when PMs also care about actual economic consequences.

Finally, it would be interesting to model the public's evaluation of PMs as a conventionally rational equilibrium inference in a model with asymmetric information. The following is an example of such a model. The economic variable follows the process

$$x^t = x^{t-1} + \mu^t + \varepsilon^t - \varepsilon^{t-1}$$

where ε^t is *i.i.d.* Each player t comes in two possible types: H and L. The prior probability over types is *i.i.d* across all players. Each player chooses between action and inaction, denoted a_1 and a_0. If $\mu^t = 0$, $a^t = a_1$ and player t is of type H, then $\mu^{t+1} = \mu$. In any other case, $\mu^{t+1} = \mu^t$. Thus, only a high-quality PM can turn the trend coefficient from zero to μ, by implementing an active reform. The evaluator is a rational player who only observes the evolution of x and the players' actions. He rewards each player according to his expected quality, given the posterior beliefs that are induced by the players' equilibrium strategies. Analyzing this interesting game of imperfect private monitoring is left for future research.

References

[1] Alesina, A. and A. Drazen (1991). "Why are Stabilizations Delayed?" American Economic Review 81, 1170-1188.

[2] Cukierman, A. and M. Tommasi (1998). "When Does it Take Nixon to Go to China?" American Economic Review 88, 180-197.

[3] Drazen, A. (2001). Political Economy in Macroeconomics. Princeton University Press.

[4] Fernandez, R. and D. Rodrik (1991). "Resistance to Reform: Status Quo Bias in the Presence of Individual Specific Uncertainty". American Economic Review 81, 1146-1155.

[5] Fong, K. (2009). "Evaluating Skilled Experts: Optimal Scoring Rules for Surgeons". Mimeo, Stanford University.

[6] Goldacre, B. (2009). Bad Science. Harper Perennial.

[7] Kahneman, D. and A. Tversky (1973). "On the Psychology of Prediction". Psychological review 80, 237-257.

[8] Orphanides, A. (1992). "The Timing of Stabilizations". Finance and Economics Discussion Series No. 194, Washington DC: Federal Reserve Board.

[9] Rothschild, M. and J. Stiglitz (1970). "Increasing Risk: I. A Definition". Journal of Economic Theory 2(3), 225-243.

[10] Rothschild, M. and J. Stiglitz (1971). "Increasing Risk: II. Its Economic Consequences". Journal of Economic Theory 5(1), 66-84.

[11] Spiegler, R. (2006). "The Market for Quacks". Review of Economic Studies 73, 1113-1131.