THE PINHAS SAPIR CENTER FOR DEVELOPMENT
TEL AVIV UNIVERSITY

Placebo Reforms

Ran Spiegler[1]

Discussion Paper No. 3-11

May 2011

I thank Ayala Arad, Eddie Dekel, Kfir Eliaz, Erik Eyster, David Shanks, Eilon Solan and seminar participants at Tel Aviv University, Hebrew University and the LSE, for their helpful comments. Financial support from The Pinhas Sapir Center for Development and the ESRC (UK) is gratefully acknowledged

[1] Tel Aviv University and University College London. Email: r.spiegler@ucl.ac.uk. URL:http://www.tau.ac.il/~rani

Abstract

I study a dynamic model of strategic reform decisions that may affect the stochastic evolution of a publicly observed economic variable. Policy makers try to maximize their public evaluation, which follows a boundedly rational rule for attributing observed outcomes to observed actions. Specifically, the public attributes recent changes to the most recent intervention. I analyze subgame perfect equilibrium in this model for a variety of stochastic processes. In particular, when the economic variable follows a (history-dependent) linear growth trend with noise, equilibrium is essentially unique and stationary, bearing a subtle formal relation to optimal search models. In equilibrium, policy makers tend to act during temporary crises, display risk aversion conditional on acting, and prefer that the random shocks associated with reforms be permanent rather than transient.

1 Introduction

I present a model of strategic policy making, in which policy makers (PMs henceforth) care about the perceived outcomes of their actions, and public perception is based on a boundedly rational rule for attributing outcomes to actions. I use this model to generate theoretical insights into aspects of PMs' reform decisions, such as the timing of reforms and the risk attitudes exhibited by choices between various policies.

To motivate our discussion, consider the following hypothetical scenario. You have been appointed as Chief of Police in a certain district. You want the public to remember you as someone who brought down crime levels. 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.

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 of their own economic reforms. How do such concerns affect decision makers' actions, particularly when they realize that their successors will face a similar dilemma?

To address this question, I construct a stylized dynamic model of strategic reform choices, in which an infinite sequence of PMs monitor the stochastic evolution of an economic variable x. Each PM moves once, and chooses an action (from a possibly history-dependent feasible set) that may or may not affect the continuation of the process that governs x. There is one action, denoted 0 and interpreted as a "default" or "inaction" option, which is always feasible. I will often refer to all other actions as "active reforms" or "interventions".

The sole objective of each PM is to maximize the outcome of his evaluation by the public. Specifically, PMs would like the public to attribute good outcomes to their own actions and bad outcomes to other PMs' actions. Public evaluation takes place at any period, and each PM employs a constant discount factor δ to weigh all future evaluation periods. A PM with a large δ cares mostly about evaluations in the distant future - his "legacy" - while a PM with a small δ has short-term career concerns and therefore cares mostly about proximate evaluations.

The public's attribution rule is a crucial component of the model. Here I introduce a modeling innovation and assume that the rule departs from conventional "rational expectations". The motivation is that in the class of situations I am interested in - from evaluation of sports managers by fans to evaluation of politicians by voters - it makes sense to assume that evaluators lack the decision makers' degree of sophistication. This is not an informational asymmetry in the usual sense, but rather an asymmetry in the quality of understanding of the underlying stochastic model. In these settings, assuming that non-experts rely on an intuitive heuristic for drawing links between actions and outcomes has considerable appeal.

Of course, there is a variety of boundedly rational attribution rules that one could assume. I impose the following: *changes in x are always attributed to the most recent intervention*. That is, at any period t, the public considers the latest period $s < t$ in which a PM chose an action $a \neq 0$, and attributes the entire difference $x^t - x^s$ to the PM who moved at period s.

This attribution rule captures a common intuition about causality: events that are both *salient* and *recent* are intuitively perceived to be causes of an observed outcome.[1] An active intervention is intuitively more salient as a failure to act, and therefore more likely to be perceived as a cause. Thus, for example, when a patient's medical condition improves, we tend to attribute the recovery to the latest medical treatment he received. Similarly, when a sports team's performance improves shortly after its manager has been replaced, fans tend to attribute the recovery to the change. Finally, in a bargaining situation, when one concession immediately follows another after a long stalemate, we tend to guess that there is a causal link between the two concessions. I basically assume that the public attributes changes in x to PMs' actions along similar lines.

Because the public's attribution rule is not based on a thorough understanding of the underlying stochastic process, it can generate systematic errors, such as giving a PM credit for a recovery that was purely due to chance. Psychologists (notably Kahneman and Tversky (1973)) have demonstrated that when people identify an intuitive casual link between an observed outcome and a preceding event, they display a strong tendency to embrace it, even when it overrides sound statistical reasoning. In our context, intuitive attribution of recent changes in the economic variable to recent interventions can be fallacious, partly because it ignores the fact that the timing of the intervention is endogenous and reflects a selection bias.

[1] For psychological research on intuitive causality judgments, see Shanks et al. (1989), Sloman (2005) and Lagnado and Speekenbrink (2010).

For example, 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 not entirely 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*" (to be distinguished from a truly physiological effect; it may be exhibited by other observers than the patient himself). In fact, it is often argued (see Goldacre (2009, pp. 38-39)) that pressure from patients guided by this type of naive inference is one of the factors that have led to the growing abuse of antibiotics.

A similar statistical Placebo effect is at play in the context of our model, where naive public inferences distort PMs' incentives as they contemplate their reform strategies. In particular, when the stochastic process is mean reverting, PMs have an incentive to implement an active reform following a negative shock, even when the reform has no real impact on economic performance, anticipating that public evaluation will neglect the mean reversion. I refer to interventions that do not affect the continuation of x, and whose sole purpose is to take advantage of the public's boundedly rational attribution rule, as "*Placebo reforms*".

This observation has interesting implications for the PMs' strategic considerations. The attribution rule implies that if the PM who moves at period t chooses to intervene, he will not get credit for any changes in x that take place after the next intervention. This means that the PM will never get credit for developments that follow the next intervention. However, we have just observed that the next intervention is endogenous and exhibits "adverse selection", in the sense that it tends to follow negative shocks. Therefore, the expected credit that the PM will get for changes in x is lower than if he did not face any successors. In other words, there is a "strategic multiplier" of the incentive to choose the default.

An illustrative example

Suppose that the value of x evolves according to an entirely deterministic cyclical process that is independent of the PMs actions. All that PMs choose is whether or not to act. For simplicity, assume that $x^t \neq x^s$ for every two periods t, s that belong to the same cycle. Then, if the PMs' discount factor is sufficiently close to one, subgame perfect equilibrium has a simple structure: every player t chooses to intervene if and only if x^t attains the *minimal* level in the cycle.

The reasoning behind this result is simple. First, player t acts whenever x^t hits the minimal level, because regardless of the future PMs' strategies, the average change in the value of x is strictly positive. Now, let x^* be the maximal value of x for which PMs sometimes choose to act, and suppose that x^* is above the minimal level. Since we have already established that PMs act at least once per cycle, the PM's payoff is primarily determined by the time of the next intervention. This means that if player t acts when $x^t = x^*$, the expected change in the value of x until the time of the player's evaluation is negative, contradicting the assumption that the player chooses to act.

This result crucially relies on the strategic aspect of the model, namely the endogeneity of the timing of the next intervention. Consider an alternative model, in which player t is the last PM to move, such that all future developments are attributed to the player if he intervenes. Then, when the discount factor is sufficiently close to one, player t will act whenever x^t falls below the *average* value of x over the cycle. Thus, the PM's realization that his payoff will be determined by the strategic considerations of future PMs implies a significantly weaker incentive to act.

The reasoning in this example is reminiscent of "unraveling" arguments in adverse-selection models (e.g., Milgrom (1981)). And indeed, the model of this paper introduces a novel adverse selection effect. When a PM chooses whether to implement an active reform, he realizes that subsequent PMs tend to act following relatively low realizations of x. As a result, the value of x conditional on the event of a future intervention is biased downward (relative to the unconditional discounted expected value of x). This means that the PM's payoff from acting suffers a downward bias, and this effect is analogous to adverse-selection effects in models of trade under asymmetric information. Of course, the adverse selection in the present model is entirely endogenous and has nothing to do with asymmetric information. Nevertheless, the adverse selection analogy is useful for understanding how the model works.

Overview of the results

After presenting the model in Section 2, I turn to equilibrium characterizations for classes of stochastic processes that are of interest from the point of view of economic applications.

In Section 3, I assume that x follows a growth process with a linear trend and independently distributed noise. Both the trend slope and noise distribution are determined by the most recent active reform decision. Subgame perfect equilibrium turns out to be subtly related to stationary search models. In equilibrium, each PM chooses to intervene if and only if the noise realization drops below a unique, stationary cutoff. Conditional on implementing a reform, the PM chooses from the set of actions that maximizes a very simple target function that exhibits risk aversion, as it trades off the expected return from an action and the riskiness of its noise distribution. When this set is a singleton, subgame perfect equilibrium is unique and stationary, and all interventions along the equilibrium path, except possibly the first one, are Placebo reforms. When the noise associated with each action may have a permanent component, the PMs' equilibrium behavior displays a taste for permanent shocks. Both this "permanence seeking" and the risk aversion highlighted above are features that crucially rely on the strategic nature of the model; they disappear in a model with a single PM who acts once and faces no successors.

In Section 4, I analyze a simple example of crisis dynamics, and show how PMs' career/legacy concerns can exacerbate a deterioration and make it irreversible. In Section 5, I focus on environments in which player 0's action determines an irreversible stochastic process, such that all subsequent interventions are by definition Placebo reforms. I show that subgame perfect equilibrium payoffs are unique and possess a simple recursive characterization, which I then employ to highlight the intertemporal trade-offs that player 0 face when contemplating reforms that induce growth processes with $i.i.d$ noise and a time-varying trend. I show that the PM's equilibrium behavior displays a non-trivial, partial form of myopia. The concluding section is devoted to a brief discussion of variations and extensions of the model as well as related literature.

2 A Model

An economic variable x evolves over (discrete) time, $t = 0, 1, 2, ...$ according to some stochastic process. In each period t, a distinct PM, referred to as player t, observes the entire history $h = (x^0, a^0, ..., x^{t-1}, a^{t-1}, x^t)$, where x^s and a^s denote the realization of x and the action taken at period s, respectively. The process is determined entirely by this sequence of actions and outcomes: no additional variables that are hidden from PMs are relevant for its continuation. I often use $t(h)$ to denote the identity of the player who moves at h, and $x(h)$ to denote $x^{t(h)}$. Upon observing the history, player t chooses an action a^t from a set of feasible actions $A(h)$. The action may affect the continuation of the stochastic process. Assume that $|A(h)| \geq 2$ for every history h, and that there exists a "null action" 0, interpreted as *inaction* or as a *default*, such that $0 \in A(h)$ for every h. Any $a \neq 0$ is interpreted as an active reform strategy. The stochastic process is common knowledge among all PMs.

To complete this description into a full-fledged infinite-horizon game with perfect information, we need to describe the players' preferences. Along a given path of the game, for any period t, define $r(t)$ as the *earliest* period $r > t$ in which $a^r \neq 0$. If none exists, then $r(t) = \infty$. Player t's payoff is

$$\begin{cases} (1-\delta)\sum_{s=t+1}^{\infty} \delta^{s-t-1} x^{\min[s,r(t)]} - x^t & if \quad a^t \neq 0 \\ 0 & if \quad a^t = 0 \end{cases}$$

where $\delta \in (0,1)$ is a discount factor. Note that when δ tends to one, player t's payoff from playing $a \neq 0$ converges to $x^{r(t)} - x^t$. Throughout the paper, I take it for granted that when a PM is indifferent between active reform and the default, he goes for the latter.

The interpretation of this payoff function is as follows. When a PM remains inactive, none of the changes in the economic variable are attributed to him, because they are all attributed to other PMs' interventions. 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 are attributed to him. The discount factor captures the PM's horizon. When δ is close to zero, the PM is motivated by short-term career concerns: he cares about how the public will evaluate him in the short run. When δ is close to one, he cares about how his actions will be regarded in the eyes of posterity.

As in Example 3.1, we see that the PMs' anticipation of future PMs' career/legacy motive is a "strategic multiplier" that exacerbates the tendency not to intervene. In Example 3.1, this had no effect on the evolution of the economic variable because by assumption all interventions were "Placebo reforms". In contrast, in Example 3.2, the disincentive to act implies that the system will permanently remain in a state of deep crisis.

3 Stationary Growth Processes

I now turn to the main application of the model, where the stochastic process follows a linear growth trend with independently distributed, transient noise, such that both the trend and the noise distribution are determined by the latest intervention. Formally, assume that the set of feasible actions is fixed throughout the game, and denoted A. (For expositional simplicity, I state the model and the results for finite 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 active reform implemented prior to t. Formally, let t' be the latest period $s < t$ for which $a^s \neq 0$; then, $b^t = a^{t'}$. If no such period t' exists, set $b^t = 0$. The economic variable x evolves according to the following equation:

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

where $\mu^t = \mu_{b^t}$ and ε^t is an independent new draw from f_{b^t}. Note that since every player t perfectly observes the entire history of actions and realizations of x, he knows the realizations ε^t and ε^{t-1}.

3.1 The Riskiness Function

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

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

for every $\varepsilon \in [-k, k]$. 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:

$$R(\varepsilon) \equiv \varepsilon \cdot F(\varepsilon) - \int_{-k}^{\varepsilon} zf(z)dz \tag{2}$$

Finally, it is easy to check that: *(i)* the function $R(\varepsilon) - \varepsilon$ is non-negative and strictly decreasing with ε; *(ii)* $R(\varepsilon) - \varepsilon \leq \frac{1}{2}(k - \varepsilon)$ for all ε (this weak inequality is binding at $\varepsilon = k$, because $R(k) = k$).

3.2 Subgame Perfect Equilibrium

For every $a \neq 0$, 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 1 *In any subgame perfect equilibrium, each player t chooses 0 whenever $\varepsilon^t > \varepsilon^*$, and an action in $\arg\max_{a\neq 0} \Pi(a, \varepsilon^*)$ whenever $\varepsilon^t < \varepsilon^*$, where ε^* is uniquely defined by the equation*

$$\max_{a\neq 0} \quad \Pi(a, \varepsilon^*) = (1-\delta)\cdot\varepsilon^* \tag{3}$$

Proof. The proof is structured as follows. I first establish lower and upper bounds on the equilibrium payoff that each player t can attain from choosing an action $a \neq 0$. These bounds do not differ across players because for every player t, the only aspect of the history that is relevant for the set of feasible payoffs is ε^t. I next show that the two bounds coincide, and use this to pin down the equilibria.

In what follows, I define player t's *gross continuation payoff* from choosing $a \neq 0$ to be equal to his payoff from this action plus ε^t.

Step 1: A lower bound on gross continuation payoff

Proof: Consider an arbitrary finite history h. Let $C(h, a)$ denote the set of finite histories in the subgame that begins after player $t(h)$ chooses the action a. Denote $t(h) = t$. To obtain a lower bound on player t's gross continuation payoff, we need to find a strategy profile in this subgame that minimizes the expectation of

$$(1-\delta)\sum_{s=t+1}^{\infty}\delta^{s-t-1}\cdot\left[\varepsilon^{\min(s,r(t))} + (\min(s, r(t)) - t)\cdot\mu_a\right]$$

This is equivalent to finding a stopping rule $\alpha : C(h, a) \to \{\mathit{stop}, \mathit{continue}\}$ - defined by $\alpha(h') = \mathit{continue}$ if and only if player $t(h')$ plays 0 - that solves a stationary stopping problem of searching for a low price, where μ_a is the constant cost of search per period; ε^s is the price encountered in period s, drawn *i.i.d* according to the density function f_a; and $1-\delta$ is a constant exogenous stopping probability.

In this well-known textbook problem (e.g., see Stokey, Lucas and Prescott (1989, pp. 304-315), the optimal stopping rule follows a stationary cutoff: stop in period s if and only if $\varepsilon^s < \varepsilon^*_a$, where ε^*_a is uniquely given by the equation

$$\mu_a = \delta\cdot\int_{-k_a}^{\varepsilon^*_a}(\varepsilon^*_a - \varepsilon)f_a(\varepsilon) + (1-\delta)\cdot\varepsilon^*_a$$

which can be rewritten as

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

Moreover, the gross continuation payoff induced by this stopping rule is ε^*_a. Therefore, the minimal payoff that player t can secure from implementing an active reform is $\max_{a\neq 0} \varepsilon_a^* = \varepsilon^*$.

Step 2: An upper bound on gross continuation payoffs

Proof: Let V^* denote the highest gross continuation payoff that any player can attain in equilibrium conditional on choosing $a \neq 0$. Then:

$$V^* \leq \sup_{a\neq 0} \left\{ \mu_a + (1-\delta)\cdot \int_{-k_a}^{k_a} \varepsilon f_a(\varepsilon) + \delta \cdot [\int_{-k_a}^{\varepsilon^*} \varepsilon f_a(\varepsilon) + (1 - F_a(\varepsilon^*))\cdot V^*] \right\}$$

The reason is as follows. First, by Step 1, in any continuation, player $t+1$ will choose $a \neq 0$ whenever $\varepsilon^{t+1} < \varepsilon^*$. This explains the first term in the square brackets. Second, the discounted sum of payoff flows that accrue to player t in periods $s \geq t+1$ conditional on $a^{t+1} = 0$ is by definition bounded from above by V^*. Rearranging this inequality, and exploiting the fact that the expected value of ε is zero, we obtain

$$V^* \leq \varepsilon^*$$

Thus, the upper and lower bounds on gross continuation payoffs coincide. Therefore, equilibrium gross continuation payoffs are uniquely given by ε^*.

Step 3: Pinning down equilibria

Proof: By the previous step, whenever player t chooses $a \neq 0$ in equilibrium, he necessarily chooses $a \in \arg\max_{a\neq 0} \varepsilon_a^*$, and he chooses $a = 0$ if and only if $\varepsilon^t \geq \varepsilon^*$. Note that by definition:

$$\begin{aligned}
\varepsilon^* &= \max_{a\neq 0} \quad \mu_a + \delta \cdot [\int_{-k_a}^{\varepsilon^*} \varepsilon f_a(\varepsilon) + (1 - F_a(\varepsilon^*))\cdot \varepsilon^*] \\
&= \max_{a\neq 0} \quad [\mu_a - \delta R_a(\varepsilon^*) + \delta \cdot \varepsilon^*] \\
&= \max_{a\neq 0} \Pi(a, \varepsilon^*) + \delta \cdot \varepsilon^*
\end{aligned}$$

This completes the characterization of equilibrium. ■

Thus, independently of the history, each player t chooses to intervene if and only if the noise realization in period t is below a stationary cutoff ε^*. Conditional on intervening, he chooses an action $a \neq 0$ that maximizes $\Pi(a, \varepsilon^*)$. If $\arg\max_{a\neq 0} \Pi(a, \varepsilon)$ is unique for all ε, then the equilibrium is necessarily unique, such that each player t chooses $a^* = \arg\max_{a\neq 0} \Pi(a, \varepsilon^*)$ whenever $\varepsilon^t < \varepsilon^*$. In this case, only the first PM who plays $a \neq 0$ brings a real change in the expected trend, from zero to μ_{a^*}. From that moment, the expected trend is μ_{a^*} forever, and all subsequent interventions along the equilibrium path are "placebo reforms".

Equilibrium properties

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

Timing. The equilibrium timing of active reform follows a stationary cut-off rule: each player t chooses $a \neq 0$ 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 models

The proof of Proposition 1 makes use of a formal analogy to a textbook search problem. Let us explore the intuition behind this analogy. First, suppose that $A = \{0, 1\}$ and take the $\delta \to 1$ limit. In equilibrium, each player t chooses $a^t = 1$ 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 cutoff rule in a textbook optimal stopping problem, in which a consumer, say, searches sequentially for a low price drawn from a stationary distribution 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.

However, the analogy comes with a twist, because the meaning of the actions 0 and 1 is not stable over time: for the current player, 0 means stopping, while his calculation of the optimal action implies that for all subsequent players, 0 means continuing. Thus, in equilibrium PMs behave as if they collectively solve a textbook stopping problem of searching for a low price, except that their stopping decision with respect to the optimal cutoff is inverted. When $|A| > 2$, the "inverted search" analogy is extended to allow for multiple search pools, as if each action $a \neq 0$ gives the consumer access to a different search pool with a characteristic (stationary) price distribution. The equilibrium risk aversion displayed by PMs in our model is thus a mirror image of the preference for high-variance price distributions exhibited by optimal behavior in the analogous search model.

The effect of a longer evaluation horizon
How is equilibrium behavior affected by changes in the PMs' evaluation horizon, as captured by the discount factor δ?

Corollary 1 *When $A = \{0, 1\}$, the equilibrium cutoff ε^* decreases with δ.*

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

$$\mu_1 = \delta \cdot R(\varepsilon^*) + (1 - \delta) \cdot \varepsilon^* \tag{4}$$

In Section 3.1, we saw that $R(\varepsilon) > \varepsilon$ for all $\varepsilon < k$. Therefore, in order to satisfy (4), ε^* must decrease when δ goes up.

Thus, the more PMs care about "posterity" (that is, evaluations that lie in the distant future), the more reluctant to intervene they become, such that the equilibrium probability of reform goes down. The intuition is that a longer horizon increases the weight of future PMs' intervention decisions in the calculation of the current PM's payoff. Because of the adverse selection that characterizes such a decision, the current PM's disincentive to act becomes stronger. The formal link to optimal stopping models 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 contains more than two actions, the effects of a longer horizon on the equilibrium timing of reform and PMs' risk attitudes become more subtly intertwined. It is easy to see that holding the cutoff ε^* fixed, PMs become more risk averse as δ goes up. However, since ε^* is endogenously determined, it appears that stronger assumptions are required to obtain clear-cut results.

For example, let $A = [l, h] \cup \{0\}$, 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 $r \in (0, 1)$ is an exogenous constant. As to the default action, all we need to assume is that the noise distribution associated with it has a sufficiently large support (this assumption is relevant only for histories in which all prior PMs chose the default). To characterize subgame perfect equilibria, we need to find a noise realization ε^* such that:

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

The solution to this problem induces a reform probability of

$$\frac{-(1-\delta) + \sqrt{1-\delta(1-r)}}{\delta} \tag{5}$$

It can be verified that this expression increases with r and decreases with δ. In the $\delta \to 1$ limit, reform probability is $\sqrt{r}$.

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, as long as $r < \frac{1}{4}$, PMs opt for the lowest-risk, lowest-return (highest-risk, highest-return) reform strategy when the discount factor is high (low).

Ex-ante payoffs for patient PMs

Suppose that just before player t observes ε^t, he is asked to evaluate his equilibrium expected payoff. In the $\delta \to 1$ limit, we obtain

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

where $a^* \in \arg\max_{a\neq 0} \Pi(a, \varepsilon^*)$. It is easy to see that this expression is equal to $R_{a^*}(\varepsilon)$, and therefore, by Proposition 1, to μ_{a^*}. That is, the player's ex-ante expected payoff is equal to the trend parameter that characterizes the active reform strategy taken in equilibrium.

This observation has interesting welfare implications. On one hand, it is plausible to assume that the public's welfare criterion is to maximize long-run growth. On the other hand, PMs turn out to use this same criterion to evaluate equilibrium outcomes. However, the equilibrium action a^* does not maximize long-run growth, but a target function that trades off the growth rate associated with reform strategies and their riskiness. In this sense, the equilibrium outcome is inefficient: if the public had a correct understanding of the model, all parties would agree that the equilibrium path is sub-optimal because it displays insufficient growth.

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 any subsequent PM to act. The PM's payoff function is exactly as in the model presented in Section 2, except that $r(t) = \infty$ with certainty.

The PM's expected payoff from taking an action $a \neq 0$ in period t, given x^t and ε^t, is

$$\frac{\mu_a}{1-\delta} - \varepsilon^t \tag{6}$$

Thus, player t will act if and only if $\varepsilon^t \leq \tilde{\varepsilon}$, where $\tilde{\varepsilon}$ is given by

$$\max_{a \neq 0} \mu_a = (1-\delta)\tilde{\varepsilon} \tag{7}$$

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

Compare this with the PMs' equilibrium behavior in our model, given by (3). 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 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 component 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 determines 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.

3.3 Permanent Shocks

In this sub-section, I extend the growth model of this section by incorporating the possibility of permanent shocks. Specifically, assume that every action a is characterized by an additional parameter $\rho_a \in [0, 1]$, and that the process (1) is modified as follows:

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

where $\rho^t = \rho_{b^t}$. Thus, ρ^t is the probability that the random, independent shock in period t is permanent; and just like the trend parameter and noise distribution, ρ^t is determined by the latest intervention prior to t.

The characterization of subgame perfect equilibrium undergoes a slight modification under this extension. For every $a \neq 0$, define

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

As before, to ensure an interior solution, assume $\mu_a^* \in (0, k_a)$ for every $a \neq 0$.

Proposition 2 *Subgame perfect equilibrium is characterized exactly as in Proposition 1, except that μ_a^* substitutes μ_a.*

I omit the proof, as it follows the same outline as the proof of Proposition 1. It immediately follows from this characterization that in equilibrium, all PMs display a preference for reform strategies that are associated with permanent shocks. Moreover, PMs are less risk averse when they choose actions with more permanent shocks. The intuition is simple. Suppose that reform involves appointing a new administrator. The impact of any new administrator on the economic variable is a function of his personal characteristics, and therefore uncertain. However, it is likely to be durable if the administrator is appointed for a long term. This attenuates the mean reversion that causes the adverse selection effect on subsequent PMs' timing decisions, and therefore the incentive to act and take risks goes up.

4 Crisis Dynamics

There is a common intuition that when an economic system enters a state of crisis, attempts by PMs to fix it tend to arrive too late, and in the meantime the crisis continues to spiral. In this section I examine a simple example that demonstrates how this tendency can be exacerbated by the career/legacy concerns captured by our model.

Assume that x gets values in $\{-N, ..., -1, 0\}$, $N \geq 2$. A value $x < 0$ represents a state of crisis; the lower the value of x, the deeper the crisis. Assume that $x^0 < 0$. Each PM faces the same set of actions $A = \{0, 1\}$. The value $x = 0$ represents full recovery from the crisis. This state is absorbing: $x^t = 0$ implies $x^{t+1} = 0$, independently of player t's action. If $x^t < 0$, then $x^{t+1} = x^t + 1$ with probability $p(x^t) \cdot a^t$, and $x^{t+1} = \max(x^t - 1, -N)$ with the remaining probability, where $p(x) > 0$ for every $x < 0$.

In a non-strategic model in which a single PM makes all the decisions, it is clear that in the $\delta \to 1$ limit, the PM would always choose $a = 1$. Another relevant benchmark is the case of fully myopic PMs ($\delta = 0$). In this case, when $-N < x^t < 0$, player t would choose $a = 1$ if and only if $p(x^t) > \frac{1}{2}$. Let us turn back to our model, focusing on the $\delta \to 1$ limit. When $x^t = -N$, $a = 1$ is a dominant strategy for player t. The question is how players would act in equilibrium at other states. It turns out that under mild upper bounds on the values of $p(x)$, the existence of a single level of x at which the default is myopically optimal suffices to "infect" the players' strategic reasoning elsewhere in the game such that all PMs choose the default in equilibrium, except when the economic variable hits "rock bottom".

Proposition 3 *Suppose that $p(x) < \frac{N+x}{N+x+1}$ for all $x \in \{-N+2, ..., -1\}$. If $p(\tilde{x}) < \frac{1}{2}$ for some $\tilde{x} \in \{-N+1, ..., -1\}$, there is a unique subgame perfect equilibrium, in which each player t chooses $a = 0$ whenever $x^t > -N$.*

Proof. Clearly, $a^t = 0$ whenever $x^t = 0$, since $x^s - x^t = 0$ for every $s > t$. From now on, consider only histories h for which $x(h) < 0$. The proof proceeds stepwise.

Step 1: Players choose $a = 0$ at $\tilde{x}$.

Proof: First, suppose that $x(h) = \tilde{x}$. If player $t(h)$ chooses $a = 0$, his payoff is zero. Suppose that he chooses $a = 1$. If $x^{t(h)+1} = \tilde{x} - 1$, this implies that $x^{r(t(h))} \leq \tilde{x} - 1$. If $x^{t(h)+1} = \tilde{x} + 1$, this implies that $x^{r(t(h))} \leq \tilde{x} + 1$. Therefore, player $t(h)$ earns a payoff that is bounded from above by $p(\tilde{x}) \cdot 1 - (1 - p(\tilde{x})) \cdot 1 < 0$. Therefore, $a = 1$ is strictly dominated at h.

Step 2: Players choose $a = 0$ at every $x \in \{-N + 1, ..., \tilde{x} - 1\}$.

Proof: Consider a history h for which $x(h) \in \{-N + 1, ..., -1\}$ and it is known that player $t(h) + 1$ chooses $a = 0$ at the history $(h, 1, x(h) + 1)$. (By Step 1, such a history must exist if $\tilde{x} > -N + 1$.) Suppose that player $t(h)$ chooses $a = 1$. If $x^{t(h)+1} = x(h) + 1$, then $x^{r(t(h))} \leq x(h)$. If $x^{t(h)+1} = x(h) - 1$, then $x^{r(t(h))} \leq x(h) - 1$. It follows that the player's expected payoff is negative, hence it is optimal for player $t(h)$ to choose $a = 0$.

Step 3: Players choose $a = 0$ at every $x \in \{\tilde{x}, ..., -1\}$.

Proof: I prove this claim by induction. By Step 1, it holds for $x = \tilde{x}$. Suppose that the claim holds for some $x \in \{\tilde{x}, ..., -2\}$, and consider a history h such that $x(h) = x + 1$. If $x^{t(h)+1} = x(h) - 1$, then by the inductive step as well as Step 2, all subsequent players choose $a = 0$ until x reaches the minimal value $-N$, in which a PM will intervene. Therefore, $x^{r(t(h))} = -N$. On the other hand, if $x^{t(h)+1} = x(h) + 1$, then $x^{r(t(h))} \leq x(h) + 1$. It follows that player $t(h)$ earns an expected payoff that is bounded from above by $p(x + 1) \cdot 1 - (1 - p(x + 1)) \cdot (x + 1 + N)$. This expression is negative by the assumption that $p(x + 1) < \frac{N+x+1}{N+x+2}$. Therefore, it is optimal for player $t(h)$ to choose $a = 0$. This completes the proof. ■

Thus, along the equilibrium path, the system reaches the bottom states $-N$ and $-N + 1$ in the shortest time possible, and then fluctuates between them indefinitely. Note that this conclusion holds even in the extreme case in which N is very large, $x^0 = -1$ and $p(x^0)$ is close to $1 - \frac{1}{N}$, such the myopic behavior would lead to recovery almost with certainty. We can see that relative to the myopic benchmark, the career/legacy motives captured by our model induces a very strong disincentive to act.

The proof of Proposition 3 follows a "contagion" argument. The rough intuition is as follows. When intervention is myopically sub-optimal at a certain level of x, it is never chosen in equilibrium (except when $x = -N$). This "infects" the strategic considerations of all other players at all other interior levels of x, because when a player anticipates that his immediate successor will choose the default, he knows that the system will deteriorate as a result, and that the public will attribute this deterioration to his own intervention.

Comment on the public's attribution rule in the context of crisis dynamics

By assumption, the public's attribution rule does not give PMs who choose the default any credit (positive or negative) for subsequent developments. This aspect of the model is quite problematic in the context of the current model. In the stationary growth model of Section 3, default meant inertia. Therefore, it made some sense to give no credit to a PM who does not intervene. In contrast, in the present model, the choosing the default option results in immediate deterioration. If the public understood that, it would certainly penalize PMs who fail to act. It follows that the fallacy embodied in the public's attribution rule is greater - and therefore perhaps less plausible - in the present model, compared with the model of Section 3.

5 Irreversible Reforms

In this section, I turn to environments in which player 0's action induces an irreversible Markov process, such that his actions alone affect the evolution of x, whereas subsequent PMs' interventions are "Placebo reforms". Formally, assume that for every $t > 0$:

$$x^t = x^{t-1} + d^t \tag{8}$$

where the initial condition is $x^0 = 0$, and d^t is governed by the following stochastic process. Let Q be a finite partition of the set of finite game histories. I typically refer to elements of Q as *states*. Let $q^t \in Q$ denote the state of the process at time t. Whenever $q^t \neq q^0$, the set of feasible actions for player t is $\{0, 1\}$. For every $q \neq q^0$, $d(q)$ represents the change in the value of the economic variable when the process is in the state q - i.e., $d^t = d(q^t)$. Let A^0 denote the set of actions available to player 0. For every $q \neq q^0$, let $\tau(q' \mid q)$ be the probability that the process switches to the state $q^{t+1} = q'$ conditional on $q^t = q$. Assume $\tau(q^0 \mid q) = 0$ for every $q \neq q^0$. For every $a \in A^0$, let $\tau^0(q \mid a)$ denote the probability that the system switches from the initial state q^0 to the state $q^1 = q$ conditional on $a^0 = a$. Assume that $\tau^0(q^0 \mid a) = 0$ for every $a \in A^0$. The notion that player 0's actions are irreversible is captured by the assumptions that q^0 is never visited again after period 0 and that transitions do not depend on the actions of players $t > 0$.[2]

It turns out that in this environment, subgame perfect equilibrium payoffs are unique and given by a simple recursive characterization.

Proposition 4 *In subgame perfect equilibrium:*
(i) Each player $t > 0$ chooses $a = 1$ and earns a payoff of $V(q^t)$ if and only if $V(q^t) > 0$,

[2] The assumption that the output of states is the *change* in x (rather than, say, the level of x) is an arbitrary modeling choice suited for the application I analyze later in this section.

where V is uniquely determined by the following recursive equation:

$$V(q) = \sum_{q'} \quad [d(q') + \delta \cdot \min(0, V(q'))] \cdot \tau(q' \mid q) \tag{9}$$

(ii) Player 0 *chooses an action* $a^0 \in \arg\max_{a \in A^0 \setminus \{0\}} \tilde{V}(q^0, a)$ *if* $\tilde{V}(q^0, a^0) > 0$*, and* $a = 0$ *if* $\max_{a \in A^0 \setminus \{0\}} \tilde{V}(q^0, a) \leq 0$*, where*

$$\tilde{V}(q^0, a) = \sum_{q \neq q^0} \quad [d(q) + \delta \cdot \min(0, V(q))] \cdot \tau^0(q \mid q^0, a) \tag{10}$$

Proof. Fix an equilibrium strategy profile σ. Because x^t is governed by a Markov process and game histories are fully observed by players, it is legitimate to write a finite history at which some player $t > 0$ moves as a sequence of actions and states $(q^0, a^0, q^1, a^1, ..., a^{t-1}, q^t)$. For any such history h, let $q(h)$ denote $q^{t(h)}$, and let $V^*(h \mid \sigma)$ be the expected payoff that player $t(h)$ attains if he chooses $a = 1$. Note that his equilibrium payoff is by definition

$$U(h \mid \sigma) = \max(0, V^*(h \mid \sigma)) \tag{11}$$

because he can guarantee a payoff of zero by choosing the default. Observe that for every two periods $s > t$:

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

By definition, $r(t) = t + 1$ if and only if player $t + 1$ plays $a = 1$. Therefore, we can write V^* recursively as follows:

$$V^*(h \mid \sigma) = \sum_{q'} [d(q') + \delta \cdot \min(0, V^*((h, 1, q') \mid \sigma))] \cdot \tau(q' \mid q(h)) \tag{12}$$

Our objective is to verify that this recursive functional equation has a unique solution, that is moreover measurable with respect to $Q \setminus \{q^0\}$. Recall that without the minimum operator, this is a Bellman equation. The standard proof of uniqueness establishes that the value function (defined by the recursive functional equation) is a contraction mapping, using Blackwell's sufficient condition for a contraction.

Thus, all we need to do in the present context is to note that the set of finite histories is countable, and therefore trivially a metric space, as well as verify that Blackwell's sufficient conditions for a contraction hold. The proof is straightforward, and virtually identical to the case of a standard Bellman equation. Therefore, I omit the detailed proof and refer the reader to Acemoglu (2009, pp. 190-199, 544-549) for an accessible exposition of the proof.

It remains to be shown that $V^*(h \mid \sigma) = V^*(h' \mid \sigma) \equiv V(q)$ whenever $q(h) = q(h') = q$. Assume the contrary. Then, since d is only a function of q, we can permute the solutions for h and h', and this would still be a solution of (12), thus violating the uniqueness result. It follows that (9) represents the equilibrium payoff that each player t earns if he chooses $a = 1$. Expression (10) immediately follows as a best-reply for player 0. ■

The recursive function defined by (9)-(10) captures the essence of the PMs' strategic considerations in this model. When player t chooses to intervene, 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, such that player t will continue to get credit for changes in the value of x.

An application: Non-stationary growth processes

In the remainder of this section, I apply this equilibrium characterization to a variant on the stationary growth model of Section 3, in which player 0's action induces a growth process with a time-varying trend. This enables us to examine the time preferences that player 0 exhibits in equilibrium.

Formally, assume that every action $a \neq 0$ in A^0 is associated with two positive trend parameters, μ_a^1 and μ_a^2, such that

$$x^1 = \mu_a^1 + \varepsilon^1$$

and

$$x^{t+1} = x^t + \mu_a^2 + \varepsilon^{t+1} - \varepsilon^t$$

for every $t > 0$, independently of subsequent PMs' actions, where ε^t is $i.i.d$ according to a fixed noise distribution f which is symmetrically distributed over the support $[-k, k]$. Thus, μ_a^1 measures the expected short-run benefit from the action a, whereas μ_a^2 measures its expected long-run benefit.

It is straightforward to embed this description in the formalism introduced earlier in this section, except that the state space is infinite. For every $t > 0$, define the state q^t as a triple $(a^0, \varepsilon^t, \varepsilon^{t-1})$, and let $A(q) = A$ for every q, and $d(q^t) = \mu_{a^0} + \varepsilon^t - \varepsilon^{t-1}$.

The transition function is simple: for every $t > 0$, ε^t is an independent draw from f_{a^0}.

If player 0 were fully myopic, he would obviously choose $\arg\max_{a \neq 0} \mu_a^1$. In contrast, consider a non-strategic model in which player 0 does not expect any future PM to act. Then, if δ is sufficiently close to one, player 0 would choose $\arg\max_{a \neq 0} \mu_a^2$. Equilibrium behavior in our model departs from these two benchmarks.

Consider the $\delta \to 1$ limit. For every $a \in A(q^0) \backslash \{0\}$, define

$$\Phi(a) = \mu_a^1 + R^{-1}(\mu_a^2) - \mu_a^2 \tag{13}$$

and denote $\phi^* = \max_{a \neq 0} \Phi(a)$. Note that $\Phi(a)$ increases with μ_a^2, by the properties of the riskiness function described in Section 3.1.

Proposition 5 *There is a unique subgame perfect equilibrium, in which player* 0 *chooses* $a = 0$ *if* $\phi^* \leq 0$, *and an action in* $\arg\max_{a \neq 0} \Phi(a)$ *if* $\phi^* > 0$.

Proof. Fix player 0's action a. The continuation game that follows the realization of x^1 falls exactly under the class of games analyzed in Section 3, with $A = \{0, 1\}$, a constant trend parameter μ_a^2 and a fixed noise distribution f (which induces a riskiness function R). According to Proposition 1, each player $t \geq 1$ chooses $a^t = 1$ if and only if $\varepsilon^t < \varepsilon^*$, where ε^* is uniquely determined by the equation $R(\varepsilon^*) = \mu_a^2$. When player 0 takes an action $a \neq 0$, his expected payoff is

$$\begin{aligned} & \mu_a^1 + \int_{-k}^{\varepsilon^*} \varepsilon f(\varepsilon) + (1 - F(\varepsilon^*)) \cdot \frac{\mu_a^2}{F(\varepsilon^*)} \\ = \quad & \mu_a^1 - \mu_a^2 + \int_{-k}^{\varepsilon^*} \varepsilon f(\varepsilon) + \frac{\mu_a^2}{F(\varepsilon^*)} \end{aligned}$$

Since $\mu_a^2 = R(\varepsilon^*)$, we can use (2) to obtain:

$$\int_{-k}^{\varepsilon^*} \varepsilon f(\varepsilon) + \frac{\mu_a^2}{F(\varepsilon^*)} = \varepsilon^*$$

It follows that player 0's expected payoff is equal to

$$\mu_a^1 - \mu_a^2 + \varepsilon^* = \mu_a^1 - \mu_a^2 + R^{-1}(\mu_a^2) = \Phi(a)$$

Therefore, he plays $a = 0$ if $\phi^* \leq 0$, and an action that maximizes Φ otherwise. ■

To see the implications of this characterization for the intertemporal trade-offs revealed by player 0's equilibrium behavior, observe that the equilibrium probability that each player $t > 0$ chooses $a \neq 0$ is $\alpha^* = F[R^{-1}(\mu_a^2)]$. This probability increases with μ_a^2. But since F is by definition the derivative of R, it follows that the marginal rate of substitution between μ_a^1 and μ_a^2 that characterizes the function Φ is

$$\frac{\partial\Phi(a)/\partial\mu_a^1}{\partial\Phi(a)/\partial\mu_a^2} = \frac{\alpha^*}{1-\alpha^*}$$

This expression increases with μ_a^2. When μ_a^2 is close to zero, the marginal rate of substitution is close to zero, which means that player 0 is far-sighted at the margin: he prefers actions with lower short-run benefits and higher long-run benefits. As μ_a^2 goes up, the marginal rate of substitution goes up, and so player 0 becomes more myopic at the margin.

At this stage of the paper, it should come as no surprise that the driving force behind this result is the adverse selection that characterizes PMs' decision to implement an active reform. When feasible values of μ_a^2 are low, players $t > 0$ choose to act only after very low noise realizations, and this noise term largely dominates the overall change in x. Therefore, player 0 mostly cares about curbing the adverse selection effect, and this is accomplished by raising μ_a^2 even at the expense of a lower μ_a^1. In contrast, when feasible values of μ_a^2 are large, the adverse selection effect is small. Moreover, it does not take a long time on average before some player $t > 0$ plays $a \neq 0$. Therefore, the most effective thing player 0 can do to improve his evaluation is to choose a reform strategy that brings immediate benefits.

6 Concluding Remarks

My objective in this paper was to present a stylized model of strategic policy making when PMs have career/legacy concerns and they are evaluated by a public that employs a boundedly rational rule for attributing observed outcomes to observed actions. The model illuminates the much-researched subject of reform delay from a new angle, and also links it to other aspects of PMs' project selection, such as risk attitudes and intertemporal preferences. The model's very simplicity immediately suggests various extensions that were not examined here. For instance, introducing costly actions or multiple monitored economic variables would be straightforward extensions. In this concluding section, I briefly discuss interesting directions for continued research that I find less straightforward.

Multiple Equilibria

All the specifications of the model studied in Sections 3-5 gave rise to unique subgame perfect equilibrium payoffs, and often unique equilibrium strategies. This raises the question of whether equilibrium payoffs are unique in general. The following simple example demonstrates that the answer is negative.

Assume that the action set for each player is $A = \{0, 1\}$, and suppose that $x^{t+1} = x^t + 2a^t - 1$. The myopically optimal action is obviously $a = 1$. There is a subgame perfect equilibrium in which all PMs play this action. However, when δ is sufficiently large, there is another equilibrium, in which all PMs play $a = 0$ at all histories. To see why this is an equilibrium, suppose that player t deviates after some history and plays $a^t = 1$, such that $x^{t+1} = x^t + 1$. Given that all subsequent PMs adhere to their equilibrium strategy, we have $x^{t+2} = x^t$ and $x^{s+1} = x^s - 1$ for every $s = t + 3, t + 4, ...$, such that player t's discounted payoff is negative.

Note that like the environment studied in Section 3, this example is stationary. The difference lies in the effect of the default action on the evolution of x. In the model of Section 3 default meant inertia - i.e., the trend and noise distribution were unchanged. In contrast, in the present example choosing the default action affects the course of x^t. Thus, the question of multiple equilibria seems to revolve around the role of the default option. Finding conditions for uniqueness of subgame perfect equilibrium payoffs is an interesting problem that I leave for follow-up research.

Heterogeneous Discount Rates

The assumption that all PMs have the same discount factor has been made primarily as a simplifying starting point. There are various reasons for being interested in the case of heterogeneous discount rates. First, PMs of different age and at different stages of their career will have a different mixture of career and legacy concerns. For example, an old politician near the end of his career is likely to care about their "legacy" (large δ), whereas a younger politician will be motivated by short-term career concerns (low δ). An extended model that assumes a stochastic discount factor can capture this distinction. Second, our model assumed that each PM moves exactly once, thus ignoring re-election. When PMs can be re-elected, they are likely to be motivated by short-term career concerns in their first term and by legacy concerns in their final term. Thus, their own discount factor changes as the game progresses. Thorough investigation of these extensions is beyond the scope of the present paper.

Alternative Attribution Rules

This paper focused on a particular element of "boundedly rational" attribution of outcomes to actions, namely the tendency to credit the most recent intervention for changes in the monitored variable. In this section I briefly discuss alternative attribution rules.

Getting credit for "whatever happens during one's shift". The model assumes that PMs get zero credit for subsequent changes if they choose the default action, because the public attributes changes to the latest intervention. In reality, PMs do seem to get some credit for the changes during their term in office, independently of whether or not they attempted to intervene. We could capture this effect by adding a term $\gamma \cdot (x^{t+1} - x^t)$ to player t's payoff function, where $\gamma > 0$ is a constant. If γ is very large, player t will choose an action that maximizes the expectation of x^{t+1}.For intermediate values of γ, equilibrium behavior will trade-off this conventionally myopic motive with the effects highlighted in this paper.

Salience and the timing of evaluation. The model implicitly assumes that public constantly pays attention to the economic variable, such that evaluation takes place at every period. Each PM weighs all future evaluations according to his discount factor. In reality, salience of actions and outcomes plays an important role in determining public attention. For example, when a PM implements an active reform - i.e., chooses $a \neq 0$ - this is a mark of salience that invites public attention to the economic variable. An extreme way of capturing this salience effect would be to assume that evaluations take place only in periods s for which $a^s \neq 0$. Note, however, that this would be equivalent to taking the $\delta \to 1$ limit in our model. Thus, we can perfectly capture this salience effect without abandoning our model. Public attention to an economic variable can also be triggered by sharp changes in its value. We could incorporate this salience factor into our model by assuming that the probability that evaluation takes place in period s increases with $|x^s - x^{s-1}|$.

Attributing outcomes to developments prior to the most recent intervention. While the assumption that the public attributes changes in x to the latest intervention, in some scenarios it seems plausible to assume that the public will evaluate the latest intervention in light of counterfactuals that seem natural given the longer history. For example, consider a history in which all players $1, ..., s-1, s+1, ...t$ choose $a = 0$, while $a^s \neq 0$. Assume in addition that

$$\frac{x^s - x^0}{s} \ll \frac{x^t - x^s}{s - t}$$

It would be plausible for an evaluator at period t to give player s *positive* credit because his action appears to have curbed the negative trend, even if it has failed to overturn it.

Bayesian rational attribution. Finally, it would be interesting to model the PMs' evaluation by the public as the result of a conventionally rational equilibrium inference in a model with asymmetric information regarding the PMs' "type". In such a model, the set of players consists of the PMs and a rational "evaluator", who observes the history and rewards each PM according to long-run limit of the posterior probability that his type is "good". The obvious merit of this modeling strategy is that it is based on a behavioral assumption perceived to be less ad-hoc than any boundedly rational attribution rule. However, it has several, inter-related drawbacks. First, as I argued at the beginning of the paper, Bayesian rationality requires the public to possess an unrealistically thorough understanding of the underlying stochastic process. Second, the equilibrium inferences that the evaluator is required to carry out in such a model are likely to require unlimited memory. In particular, he will rely on observations in early stages of the game to update his beliefs in later stages, and he will keep updating his belief regarding early players' types. Finally, I expect the model to be intractable for most stochastic processes of interest. While the main equilibrium characterization results of this paper (Propositions 1 and 4) were obtained for general classes of stochastic processes, it is very likely that complete equilibrium characterization in the analogous Bayesian-rational model can be accomplished under special (ad-hoc?) stochastic processes. As Ellison (2006) pointed out in a different context, bounded-rationality models are often much more tractable than their Bayesian-rational counterparts, a property that enables the modeler to enrich the analysis in a number of important dimensions.

Related Literature

To my knowledge, this is the first paper to analyze theoretically public decision making when policy makers care about the way they will be evaluated by a boundedly rational audience. Although the model addresses in an abstract manner a general strategic situation and does not commit to a particular application, it is closely related to a strand in the political economics literature that deals with the question of reform timing. This literature has primarily tried to explain why socially beneficial reforms often seem to be 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 different factions as to which 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 (this is a conventional stopping problem, to be distinguished from the subtle formal analogy to search models highlighted in Section 3). For a survey of current approaches to this problem, see Drazen (2001, pp. 403-454).

There are a few precedents for the general idea of modeling interactions with/among agents who use boundedly rational attribution rules. Osborne and Rubinstein (1998) construct a game-theoretic solution concept in which each player forms an action-consequence link by naively extrapolating from a sample of observations taken from the opponents' mixed strategies. Spiegler (2004) analyzes a proto-bargaining game, in which a player's tendency to explain his opponent's concessions as the consequence of his own recent bargaining posture arises endogenously from a simplicity-based criterion for selecting equilibrium beliefs. Spiegler (2006) 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, this paper is somewhat related to the vast literature on career concerns in organizations and their implications for dynamic moral hazard situations (see Prendergast (1999) for a survey). The distorting effect of career concerns on experts' intervention decisions - particularly in the case of medical decision making - 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.

References

[1] Acemoglu, D. (2009). Introduction to Modern Economic Growth. Princeton University Press. Princeton, NJ.

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

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

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

[5] Ellison, G. (2006). "Bounded Rationality in Industrial Organization," in R. Blundell, W. Newey and T. Persson (eds.), Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress, Cambridge University Press.

[6] 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.

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

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

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

[10] Lagnado, D. and M. Speekenbrink (2010). "The Influence of Delays in Real-Time Causal Learning". The Open Psychology Journal 3, 184-195.

[11] Milgrom, P. (1981). "Good News and Bad News: Representation Theorems and Applications". Bell Journal of Economics 12, 380-91.

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

[13] Osborne M. and A. Rubinstein (1998), "Games with Procedurally Rational Players", American Economic Review 88, 834-849.

[14] Prendergast, C. (1999), "The Provision of Incentives in Firms", Journal of Economic Literature 37, 7-63.

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

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

[17] Shanks, D., S. Pearson and A. Dickinson (1989). "Temporal Contiguity and the Judgment of Causality by Human Subjects". Quarterly Journal of Experimental Psychology B 41, 139-159.

[18] Sloman, S. (2005). Causal Models; How People Think about the World and its Alternatives. Oxford University Press. New York.

[19] Spiegler, R. (2004). "Simplicity of Beliefs and Delay Tactics in a Concession Game". Games and Economic Behavior 47, 200-220.

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

[21] Stokey, N., R. Lucas and E. Prescott (1989), Recursive Methods in Economic Dynamics. Harvard University Press.