

Placebo Reforms*

Ran Spiegler[†]

January 26, 2012

Abstract

I study a dynamic model of strategic reform decisions that potentially 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 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

Imagine that you have been appointed as the 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

*This is a revised, and substantially shorter version of a 2010 working paper that appeared under the same title. I thank Ayala Arad, Eddie Dekel, Kfir Eliaz, Erik Eyster, Eeva Mauring, David Shanks, Eilon Solan, the editor and referees of this journal, and seminar participants at Tel Aviv University, Hebrew University, LSE and the Princeton Political Economy Conference, for their helpful comments. Financial support from the Sapir Centre, the ESRC (UK), and European Research Council Grant no. 230251, is gratefully acknowledged.

[†]Tel Aviv University and University College London. E-mail: r.spiegler@ucl.ac.uk. URL: <http://www.tau.ac.il/~rani>.

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's predicament is shared by many expert decision makers. A surgeon benefits when a patient attributes his recovery to an operation performed by the surgeon himself. CEOs get credit when the company's performance improves shortly after a major acquisition 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, especially 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 policy makers (PMs) monitor the stochastic evolution of an economic variable x . Each PM moves once, and chooses an action that may affect the continuation process. One action is interpreted as a "default" or "inaction", whereas all other actions are interpreted as "active reforms" or "interventions". The sole objective of each PM is to maximize his public evaluation. Specifically, PMs would like the public to attribute good (bad) outcomes to their own (other PMs') actions. Public evaluation takes place at all periods and each PM employs a constant discount factor to weigh all future evaluation periods. The PM chooses to intervene only if the discounted credit he expects to get is strictly positive.

The public's attribution rule is a crucial component of the model. I assume that the rule departs from conventional "rational expectations". The motivation for this departure is that in the class of situations I am interested in, evaluators arguably lack the experts' degree of sophistication, and often rely on intuitive heuristics for drawing links between actions and consequences. This is not an informational asymmetry in the usual sense, but rather an asymmetry in the level of understanding of a stochastic environment.

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 active reform, and attributes the entire difference $x^t - x^s$ to the PM who moved at period s . This rule captures a common intuition about causality: salient, recent events are intuitively perceived to be causes of an observed outcome. Action is more salient than inaction, and therefore more likely to be perceived as a cause. For example, when a sports team's performance improves shortly after its manager has been replaced, fans tend to attribute the recovery to the replacement. Similarly, 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.¹

Such intuitive causality judgments can generate systematic errors. Psychologists (notably Kahneman and Tversky (1973)) have demonstrated that when people identify an intuitive causal link between an observed outcome and a preceding event, they tend to embrace it, often at the expense of sound statistical reasoning. According to Banerjee and Duflo (2011, Ch. 3), false inferences of this kind are partly responsible for major distortions in the demand for health care that afflict developing countries:

“Because most diseases that prompt visits to the doctor are self-limiting (i.e., they will disappear no matter what), there is a good chance that patients will feel better after a single shot of antibiotics. This naturally encourages spurious causal associations: Even if the antibiotics did nothing to cure the ailment, it is normal to attribute any improvement to them. By contrast, it is not natural to attribute causal force to inaction: If a person with the flu goes to the doctor, and the doctor does nothing, and the patient then feels better, the patient will correctly infer that it was not the doctor who was responsible for the cure. And rather than thanking the doctor for his forbearance, the patient will be tempted to think that it was lucky that everything worked out this time but that a different doctor should be seen for future problems.”²

A similar effect is at play in our model, and it distorts 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 bad shock even when the reform has no real impact, anticipating that the public will neglect mean reversion and the intervention’s endogeneity. I refer to actions that do not affect the continuation of x and cater to the public’s boundedly rational attribution rule as “*Placebo reforms*”.

Strategic considerations magnify the incentive distortions. According to the attribution rule, a PM who intervenes does not get credit for developments that follow the next intervention. However, we have just observed that the next intervention exhibits “adverse selection”, in the sense that it tends to follow negative shocks. Therefore, the

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

²This phenomenon is widespread. Goldacre (2009, pp. 38-39) argues that pressure from patients guided by this type of inference has contributed to the growing abuse of antibiotics in the UK.

expected credit that the PM will get for changes in x is lower than if he did not face any successors. This “adverse selection” effect is a further disincentive to act.

To illustrate this point, suppose that x^t follows a deterministic cycle, independently of the PMs’ actions. Specifically, assume that $x^s = x^t$ if and only if $s - t$ is a multiple of some integer $L > 1$. All that PMs choose is whether or not to act; they are unable to affect the evolution of x . If the PMs’ discount factor is sufficiently close to one, subgame perfect equilibrium has a simple structure: every player t chooses to act if and only if x^t attains the *minimal* level in the cycle. To see why, note first that player t acts whenever x^t hits the minimum, because the expected value of x at future evaluations is by definition higher. Let x^* be the maximal value of x for which PMs sometimes act, and suppose that x^* is above the minimum. Since PMs are arbitrarily patient, their payoff is primarily determined by the value of x at the time of the next intervention. If player t acts when $x^t = x^*$, then by the definition of x^* , the expected change in x that is attributed to him is not positive, contradicting the optimality of his decision.

This result crucially relies on the endogeneity of the next intervention’s timing. By comparison, in a single-agent model, in which the PM who moves at t is the *last* PM to move, *all* future changes are attributed to him if he intervenes. If the PM is sufficiently patient, he will act whenever x^t falls below the *average* value of x .

Naturally, the precise equilibrium implications of the model depend on the stochastic process governing the evolution of x . In this paper, I assume that x follows a growth process with a linear trend and independently distributed noise. Both the trend slope and the noise distribution are determined by the most recent active reform decision. Interventions should be interpreted as reforms that may affect trends in areas such as crime, education, economic growth or the environment.

In subgame perfect equilibrium, each PM intervenes if and only if the noise realization drops below a unique, stationary cutoff. When the PM intervenes, he chooses an action that maximizes a simple, static function that exhibits risk aversion, trading off the trend slope with the riskiness of the noise distribution. When there is a unique such action, all interventions along the (unique) equilibrium path, except possibly the first one, are “Placebo reforms”. I show that this characterization is subtly related to optimal search models.

When the noise associated with each action has a permanent component - a case analyzed in Appendix A - the PMs’ equilibrium behavior displays a taste for permanent shocks. This effect, just like the equilibrium aversion to risk, relies on the strategic aspect of the model and disappears in the single-agent version of the model. Thus, one merit of the model is that it traces three behavioral phenomena, which are in principle

distinct, to the same underlying story: strategic PMs maximizing their evaluation by a boundedly rational public. Whether this linkage exists in real-life public decision making is a question for future research.

Related Literature

To my knowledge, this is the first paper to theoretically analyze public decision making when PMs care about the way they will be evaluated by a boundedly rational audience. It is related to a strand in the political-economics literature seeking 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' decisions 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. 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 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 the opponents' mixed strategies. Spiegler (2004) analyzes a proto-bargaining game, in which a player's tendency to explain his opponent's concessions as a consequence of his own recent bargaining posture arises endogenously from a simplicity-based criterion for selecting equilibrium beliefs. Spiegler (2006) models price competition in a market for a credence good, when consumers use anecdotal reasoning to evaluate the quality of each alternative.

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 expert 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, $t = 0, 1, 2, \dots$. In each period t , a distinct PM, referred to as player t , chooses an action a from a finite set A , after observing the history $(x^0, a^0, \dots, x^{t-1}, a^{t-1}, x^t)$, where x^s and a^s denote the realization of x and the action taken at period s , respectively. The action set A contains at least two elements, including a “null action” denoted by 0, which is interpreted as *inaction* or as a *default*. I interpret any $a \neq 0$ as an active reform strategy, also referred to as an “intervention”.

The players’ actions may affect the evolution of x . Specifically, assume that x^t 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. To state this formally, we need a bit of notation. Every action $a \in A$ is characterized by a *trend slope* $\mu_a > 0$ and a continuous density function f_a , which is symmetrically distributed around zero with support $[-k_a, k_a]$. I use F_a to denote the *cdf* induced by f_a . For a given play path and a given period t , let s be the latest period prior to t for which $a^s \neq 0$. If such a period s exists, define $b^t = a^s$, i.e. the most recent active reform implemented prior to t . If no such period s exists, set $b^t = 0$.

We can now formally define the stochastic process. For every $t > 0$:

$$x^t = x^{t-1} + \mu_{b^t} + \varepsilon^t - \varepsilon^{t-1} \quad (1)$$

where for every s , ε^s is an independent draw from f_{b^s} . The stochastic process is common knowledge among PMs. By the assumption that each player t fully observes past actions and realizations of x , he also knows the realizations ε^t and ε^{t-1} at the time he makes his choice.

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

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

where $\delta \in (0, 1)$ is a discount factor. Note that when δ tends to one and $r(t) < \infty$, 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, all changes in x are attributed to other PMs' interventions. If, on the other hand, the PM implements an active reform, he gets credit for the changes from that moment until another PM implements a new reform. 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 his "legacy" - namely, how posterity will regard his actions.

Suppose that $b^t \neq 0$ for some period $t > 0$. If the PM chooses an action $a^t \neq 0$ for which $(\mu_a, f_a) = (\mu_{b^t}, f_{b^t})$, the continuation of x will be as if he chose the default. However, these two options are not payoff-equivalent for the PM. In this case, we say that the intervention a^t is a "Placebo reform": it is an action with no real effect on x , which is taken solely for the purpose of claiming credit for observed changes in x .

I conclude this section by presenting two functions that will play an important role in our analysis. For any action a , define

$$R_a(\varepsilon) \equiv \int_{-\infty}^{\varepsilon} F_a(z) dz \quad (2)$$

for every $\varepsilon \in (-\infty, \infty)$. I refer to R as the *riskiness function* associated with a . This label is justified by a well-known result, due to Rothschild and Stiglitz (1970), that f_a second-order stochastically dominates f_b if and only if $R_b(\varepsilon) \geq R_a(\varepsilon)$ for every ε . Two easily verifiable properties will be useful in the sequel: (i) $R_a(\varepsilon) - \varepsilon$ is a non-negative-valued and strictly decreasing function; (ii) $R_a(\varepsilon) - \varepsilon \leq \frac{1}{2}(k_a - \varepsilon)$ for all ε (the inequality is binding at $\varepsilon = k_a$, because $R_a(k_a) = k_a$).

For any action a , define

$$\Pi(a, \varepsilon) \equiv \mu_a - \delta R_a(\varepsilon) \quad (3)$$

This function trades off the expected trend associated with an active reform strategy and its riskiness, according to weights determined by the PMs' discount factor. Define the noise realization ε^* as the (unique) solution to the following equation:

$$\max_{a \neq 0} \Pi(a, \varepsilon^*) = (1 - \delta)\varepsilon^* \quad (4)$$

To see why ε^* is uniquely defined, note that the R.H.S of this equation is continuous and strictly increasing in ε^* , while the L.H.S is the upper envelope of continuous, strictly decreasing functions, hence continuous and strictly decreasing itself. The L.H.S is higher (lower) than the R.H.S for low (high) values of ε^* , such that the two functions must

have a unique intersection. Note that if we lower the trend parameter associated with each intervention, or subject the noise distribution associated with each intervention to a mean-preserving spread, ε^* goes down.

3 Subgame Perfect Equilibrium

We are now ready for the main result of this paper: in subgame perfect equilibrium, each player t intervenes if and only if the noise realization in period t is below the cutoff ε^* ; conditional on intervening, he chooses an action $a \neq 0$ that maximizes $\Pi(a, \varepsilon^*)$.

Proposition 1 *In subgame perfect equilibrium, each player t chooses the action 0 whenever $\varepsilon^t \geq \varepsilon^*$, and an action in $\arg \max_{a \neq 0} \Pi(a, \varepsilon^*)$ whenever $\varepsilon^t < \varepsilon^*$.*

Proof. Define player t 's *gross payoff* from choosing $a \neq 0$ to be equal to his payoff from this action plus ε^t . Let us first verify that the strategy described in the statement of the result is an equilibrium strategy. Suppose that player t chooses some $a \neq 0$ and that $r(t) = t + n$. Then, player t 's gross payoff is

$$(1 - \delta) \left[\sum_{j=1}^n \delta^{j-1} j \mu_a + \sum_{j=n+1}^{\infty} \delta^{j-1} (n \mu_a + \varepsilon^{t+n}) \right]$$

Given that all players $s > t$ intervene if and only if $\varepsilon^s < \varepsilon^*$, player t 's gross payoff from choosing $a \neq 0$ is

$$\begin{aligned} & (1 - \delta) \sum_{n=1}^{\infty} F_a(\varepsilon^*) (1 - F_a(\varepsilon^*))^{n-1} \left[\sum_{j=1}^n \delta^{j-1} j \mu_a + \sum_{j=n+1}^{\infty} \delta^{j-1} \left(n \mu_a + \frac{\int_{-\infty}^{\varepsilon^*} \varepsilon f_a(\varepsilon) d\varepsilon}{F_a(\varepsilon^*)} \right) \right] \\ &= \frac{\mu_a + \delta \int_{-\infty}^{\varepsilon^*} \varepsilon f_a(\varepsilon) d\varepsilon}{1 - \delta(1 - F_a(\varepsilon^*))} \end{aligned}$$

Using integration by parts, we can redefine the riskiness function as follows:

$$R_a(\varepsilon) \equiv \int_{-\infty}^{\varepsilon} (\varepsilon - z) f_a(z) dz \quad (5)$$

such that player t 's gross payoff from choosing $a \neq 0$ can be written as

$$\frac{\mu_a - \delta R_a(\varepsilon^*) + \delta F_a(\varepsilon^*) \varepsilon^*}{1 - \delta + \delta F_a(\varepsilon^*)}$$

Since ε^* is defined by (4), the final expression is equal to ε^* when $a \in \arg \max_{a \neq 0} \Pi(a, \varepsilon^*)$, and strictly below ε^* when $a \notin \arg \max_{a \neq 0} \Pi(a, \varepsilon^*)$. Therefore, the strategy of playing $a = 0$ whenever $\varepsilon^t \geq \varepsilon^*$ and an action $a \neq 0$ that maximizes $\Pi(a, \varepsilon^*)$ whenever $\varepsilon^t < \varepsilon^*$ is a best-reply for player t .

The proof of equilibrium uniqueness proceeds by establishing lower and upper bounds on gross payoffs in equilibrium. I then show that the two bounds coincide, and use this to pin down the equilibria.

Step 1: A lower bound on gross payoffs

Proof: Let $H(a)$ denote the set of finite histories in the subgame Γ that begins after player t chooses the action a . To obtain a lower bound on player t 's gross payoff, we need to find a strategy profile in Γ that minimizes the expectation of

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

This is formally equivalent to the problem of finding a stopping rule $\alpha : H(a) \rightarrow \{\text{stop}, \text{continue}\}$ - defined by $\alpha(h) = \text{continue}$ if and only if the player who moves at 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

$$\mu_a = \delta \cdot \int_{-\infty}^{\varepsilon_a^*} (\varepsilon_a^* - \varepsilon) f_a(\varepsilon) d\varepsilon + (1 - \delta) \varepsilon_a^* \quad (6)$$

By (5), we can rewrite (6) as $\Pi(a, \varepsilon_a^*) = (1 - \delta) \varepsilon_a^*$. Moreover, the gross 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 payoffs

Proof: Let W denote the highest gross payoff that any player t can attain in equilibrium conditional on choosing $a \neq 0$. If player $t+1$ intervenes following some noise realization ε^{t+1} , player t 's gross payoff will be $\mu_a + \varepsilon^{t+1}$. By Step 1, player $t+1$ intervenes whenever $\varepsilon^{t+1} < \varepsilon^*$. If player $t+1$ chooses the default, player t 's gross payoff is weakly below $\mu_a + (1 - \delta) \varepsilon^{t+1} + \delta W$. The reason is that the stream of payoffs gathered by player t over

the periods $s > t + 1$ are calculated as if player $t + 1$ chose the action a and transferred his entire gross payoff - which, by definition, cannot exceed W - to player t . Recall that player $t + 1$ plays $a \neq 0$ if and only if his gross payoff exceeds the noise realization ε^{t+1} . Therefore, if $\varepsilon^{t+1} > \varepsilon^*$, player t 's payoff cannot be greater than $\mu_a + (1 - \delta)\varepsilon^{t+1} + \delta W$. The following inequality follows:

$$\begin{aligned}
W &\leq \sup_{a \neq 0} \left\{ \mu_a + \int_{-\infty}^{\varepsilon^*} \varepsilon f_a(\varepsilon) d\varepsilon + (1 - \delta) \int_{\varepsilon^*}^{\infty} \varepsilon f_a(\varepsilon) d\varepsilon + \delta(1 - F_a(\varepsilon^*))W \right\} \\
&= \sup_{a \neq 0} \left\{ \mu_a - \delta \int_{\varepsilon^*}^{\infty} \varepsilon f_a(\varepsilon) d\varepsilon + \int_{-\infty}^{\infty} \varepsilon f_a(\varepsilon) d\varepsilon + \delta(1 - F_a(\varepsilon^*))W \right\} \\
&= \sup_{a \neq 0} \left\{ \mu_a - \delta R_a(\varepsilon^*) + \delta F_a(\varepsilon^*)\varepsilon^* + \delta(1 - F_a(\varepsilon^*))W \right\} \\
&\leq \sup_{a \neq 0} \left\{ (1 - \delta)\varepsilon^* + \delta F_a(\varepsilon^*)\varepsilon^* + \delta(1 - F_a(\varepsilon^*))W \right\}
\end{aligned}$$

where the manipulation of the R.H.S on the third line makes use of the zero-mean property of f_a , and the final inequality on the fourth line relies on the definition of ε^* given by (4), and the redefinition of R_a given by (5). Rearranging this inequality, we obtain $W \leq \varepsilon^*$. Thus, the upper and lower bounds on gross payoffs coincide at ε^* .

Step 3: Pinning down ε^ and the equilibrium interventions*

Proof: By the previous step, whenever player t chooses $a \neq 0$ in equilibrium, he necessarily chooses an element in $\arg \max_{a \neq 0} \varepsilon_a^*$, and he chooses $a = 0$ if and only if $\varepsilon^t \geq \varepsilon^*$. This completes the proof. ■

The method of proof is to establish upper and lower bounds on players' "gross" equilibrium payoffs (i.e., ignoring the current noise realization), and use the stationarity of the game to show that the two bounds coincide (somewhat in the manner of basic bargaining theory - see Osborne and Rubinstein (1994, pp. 123-125)). The lower bound is attained via an analogy between the future PMs' (hypothetical) collective decision problem of minimizing the current PM's payoff and the textbook model of sequential search for a low price. I elaborate on this analogy below. Note that since PMs act whenever their gross payoff is above the current noise realization, the lower bound on gross equilibrium payoffs must also be the lowest noise realization at which future PMs could possibly choose the default. I use this property to derive an upper bound on his gross payoffs. Since the upper and lower bounds coincide, the cutoff ε^* can be pinned down.

Let us now list the main equilibrium properties.

Timing of interventions. The equilibrium timing of interventions follows a stationary cut-off rule: each player t chooses $a \neq 0$ if and only if the noise realization in period t drops below the cutoff ε^* . If $\varepsilon^* < k_a$, then since all noise distributions have a zero mean, the equilibrium timing of reforms displays “*adverse selection*”: the noise realization is negative on average in periods of interventions. In other words, PMs tend to implement reforms during temporary crises. This tendency is exacerbated when the environment becomes noisier. That is, if we subject the noise distribution associated with any action to a mean-preserving spread, the equilibrium cutoff ε^* will drop.³

Risk attitudes. Conditional on implementing an active reform, PMs’ choices display risk aversion. They choose a reform strategy as if they maximize a function that trades off the expected trend and the riskiness of the available reform strategies.

Placebo reforms. Suppose that for every ε , there is a unique action $a(\varepsilon)$ that maximizes $\Pi(a, \varepsilon)$. In this case, subgame perfect equilibrium is unique: each player t chooses $a(\varepsilon^*)$ whenever $\varepsilon^t < \varepsilon^*$, and $a = 0$ otherwise. Along the equilibrium path, the first PM who intervenes changes the trend and the noise distribution from (μ_0, f_0) to $(\mu_{a(\varepsilon^*)}, f_{a(\varepsilon^*)})$; all subsequent interventions are “placebo reforms” - they do not affect the evolution of x . When $\arg \max_{a \neq 0} \Pi(a, \varepsilon^*)$ is not unique, placebo reforms are a possible feature of equilibrium behavior, but not a necessary one.

Connection to optimal stopping models

Let us explore further the search analogy that the proof of Proposition 1 relies on. First, suppose that $A = \{0, 1\}$ and take the $\delta \rightarrow 1$ limit. In equilibrium, each player t chooses $a^t = 1$ if and only if $\varepsilon^t < \varepsilon^*$, where ε^* is uniquely defined by the equation $\mu_1 = R_1(\varepsilon^*)$, which can be rewritten as

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

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 and no discounting. 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 con-

³Note that if $\mu_{a^*} > k_{a^*}$, then by the property that $R_a(\varepsilon) \leq k_a$ for all a , it must be the case that $\varepsilon^* > k_{a^*}$. The adverse selection effect disappears, and PMs always intervene in equilibrium.

tinuing. 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 implies that each action $a \neq 0$ gives the consumer access to a different “search pool” with characteristic search cost and (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, conventional 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 ? When $A = \{0, 1\}$, the equilibrium cutoff ε^* decreases with δ . To see why, revisit equation (4), and suppose that $\varepsilon^* < k_1$. Recall that $R_1(\varepsilon) > \varepsilon$ for all $\varepsilon < k_1$. Therefore, when δ goes up, ε^* must decrease in order for (4) to hold. Thus, the more PMs care about “posterity”, the more reluctant to intervene they become, and 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 these future interventions, 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 timing of interventions and risk attitudes become more subtly intertwined, and 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 support of its noise distribution. In addition, assume that $\mu_a = \gamma \cdot a$ for every $a \in [l, h]$, where $\gamma \in (0, 1)$ is an exogenous constant. To characterize subgame perfect equilibrium in this case, we need to find a noise realization ε^* such that:

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

It can be shown that both ε^* and $a(\varepsilon^*)$ are decreasing in δ , as long as $\delta \neq 4\gamma$. In particular, when $\delta < 4\gamma$, $a(\varepsilon^*) = h$ and $\varepsilon^* \in (0, h)$; and when $\delta > 4\gamma$, $a(\varepsilon^*) = l$ and $\varepsilon^* \in (-l, 0)$. (When $\delta = 4\gamma$, every $a \in [l, h]$ is optimal conditional on acting, and the cutoff is $\varepsilon^* = 0$.) Thus, as PMs become more patient, the first intervention along

the equilibrium path arrives after a longer delay, and it leads to weaker growth. The frequency of subsequent Placebo reforms is also lower when δ goes up.⁴

Ex-ante payoffs for patient PMs

Consider the $\delta \rightarrow 1$ limit. Assume there is a unique action $a(\varepsilon)$ that maximizes $\Pi(a, \varepsilon)$ for every ε . Suppose that just before player t observes ε^t , he is asked to evaluate his equilibrium expected payoff. Recall from the proof of Proposition 1 that when player t acts, his gross payoff is ε^* . Therefore, his equilibrium (net) payoff given ε^t is $\max(0, \varepsilon^* - \varepsilon^t)$. It follows that the player's ex-ante payoff is

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

Thus, the player's ex-ante expected payoff is equal to the trend slope that characterizes the active reform implemented in equilibrium. This observation has interesting welfare implications. Long-run growth is not only the criterion PMs use to evaluate their ex-ante equilibrium payoffs, but also a natural criterion for evaluating the public's welfare. Therefore, the PMs and the public would appear to have common interests. However, recall that in equilibrium, PMs exhibit risk aversion and do not choose actions that maximize growth. In this sense, the equilibrium outcome is inefficient.

Comparison with a Non-Strategic Model

To get a deeper understanding of the strategic considerations in our model, it will be useful to draw a comparison with a simple single-agent model, in which player 0 does not face any successors. The PM's payoff function is exactly as in the model presented in Section 2, except that $r(0) = \infty$ with certainty. The PM's expected payoff from taking an action $a \neq 0$ is $\mu_a / (1 - \delta) - \varepsilon^0$. The PM will act if and only if $\varepsilon^0 < \tilde{\varepsilon}$, where $\tilde{\varepsilon}$ is given by $\max_{a \neq 0} \mu_a = (1 - \delta)\tilde{\varepsilon}$. Conditional on acting, he will choose a to maximize μ_a .

Compare this with the PMs' equilibrium cutoff rule given by (4). 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. First, $\varepsilon^* < \tilde{\varepsilon}$. In particular, when $\max_{a \in A} \mu_a \approx 0$, $\tilde{\varepsilon} \approx 0$ whereas in the strategic model the cutoff is bounded away below zero. The reason is that only in the strategic model, PMs are concerned with the "adverse selection" effect. Second, the structure of noise is irrelevant for the PM's decision in the single-agent

⁴Although A is not finite in this example, Proposition 1 can be conventionally extended to cover this case.

model. The reason is the same: in the absence of the adverse selection effect, the PM is risk-neutral and therefore disregards the noise structure. Finally, comparative statics with respect to δ are different. In the single-agent model, a higher δ leads to a higher reform probability (because we assumed $\mu_a > 0$ for all a). In the $\delta \rightarrow 1$ limit, the PM acts with probability one. In contrast, in the strategic model, we saw that a higher δ can lead to a lower equilibrium reform probability, because it magnifies the adverse selection effect.

4 Concluding Remarks

My objective in this paper was to present a stylized model of strategic policy making when PMs are evaluated by the public according to a boundedly rational rule for attributing observed outcomes to observed actions. The model illuminates the subject of reform delay from a new angle, and links it to other aspects of PMs' reform choice, notably risk aversion. The model's very simplicity immediately suggests various extensions that were not examined here. Some of them are trivial, e.g., assuming that interventions are costly. Others involve different stochastic processes; two examples are analyzed in Appendices A and B. Additional interesting extensions would include multiple monitored economic variables, or heterogeneous discount rates among PMs. In this concluding section, I briefly discuss two criticisms of the model.

Is it possible to “fake” a reform?

Recall that when $\arg \max_a \Pi(a, \varepsilon^*)$ is unique, all interventions along the equilibrium path (except possibly the first) are “Placebo reforms”. In particular, whenever PMs intervene they take the same action. One could argue that if the action taken by player t is manifestly identical to the one taken by the previous reformer, the public will not regard it as a *reform* and therefore will not give player t any credit for it. One way to avoid this criticism is to assume that every feasible action has infinitely many “replica” that share the same (μ, f) , yet are regarded as distinct by the public.

Alternatively, we can maintain the assumption that A is finite, and modify our model by assuming that in order for player t to get credit for an intervention, he must choose $a^t \neq 0, b^t$. This variation can generate new inefficiencies. To see why, let $A = \{0, 1, 2\}$, and assume $\mu_2 > \mu_1$, $f_1 \equiv f_2$. In the original model, the action that PMs choose conditional on intervening would always be 2. By comparison, under the alternative assumption, there is a subgame perfect equilibrium with the following structure. There are two cutoff points, $\varepsilon_1^* < \varepsilon_2^*$. As long as $b^t \neq 2$, player t acts

whenever $\varepsilon^t < \varepsilon_2^*$ and chooses $a^t = 2$. If, on the other hand, $b^t = 2$, player t acts whenever $\varepsilon^t < \varepsilon_1^*$ and chooses $a^t = 1$. Thus, when the status quo is the action 2, there will be a transition to the inferior action 1 following a sufficiently bad shock, and the transition back to the superior action 2 will occur with some delay.

Limitations of the attribution rule

Like any model of bounded rationality, the reasonableness of the attribution rule assumed in this paper varies with the context. For example, consider the assumption that PMs get no credit when they choose the default. This may seem reasonable when x^t evolves according to (1). In this case, default implies inertia; and one could argue that even by rational reckoning, a PM who fails to act should not be credited with (or blamed for) subsequent changes in x . However, for other processes, e.g., when failure to act accelerates a deterioration, this assumption may appear implausible.

The model implicitly assumes that the public constantly pays attention to the economic variable, such that evaluation takes place at every period. In reality, salience of actions and outcomes influence public attention. For example, when a PM implements an active reform, this in itself is a salient event that attracts 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 $\delta \rightarrow 1$ limit in our model.⁵ Thus, we can perfectly capture this salience effect without abandoning our model. Heightened public attention to an economic variable can also be triggered by sharp changes in its value. The attribution rule assumed in this paper cannot capture this aspect. It would be interesting to explore alternative rules that capture such wider intuitions.

A more conventional approach would be to model the PMs' public evaluation as the result of Bayesian-rational equilibrium inference in a model with asymmetric information regarding the PMs' "type", where different types have different abilities to influence the evolution of x . 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 the long-run limit of the posterior probability that his type is "good". I find this a very interesting direction for future research.

⁵The equivalence would fail if the PM expected no future interventions. However, this scenario will not materialize in equilibrium, and therefore the equivalence holds de facto.

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] Banerjee, A. and E. Duflo (2011). *Poor Economics*. PublicAffairs.
- [5] 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.
- [6] Fong, K. (2009). “Evaluating Skilled Experts: Optimal Scoring Rules for Surgeons”. Mimeo, Stanford University.
- [7] Goldacre, B. (2009). *Bad Science*. Harper Perennial.
- [8] Kahneman, D. and A. Tversky (1973). “On the Psychology of Prediction”. *Psychological review* 80, 237-257.
- [9] Lagnado, D. and M. Speekenbrink (2010). “The Influence of Delays in Real-Time Causal Learning”. *The Open Psychology Journal* 3, 184-195.
- [10] Orphanides, A. (1992). “The Timing of Stabilizations”. *Finance and Economics Discussion Series No. 194*, Washington DC: Federal Reserve Board.
- [11] Osborne M. and A. Rubinstein (1994). *A Course in Game Theory*. MIT Press.
- [12] Osborne M. and A. Rubinstein (1998), “Games with Procedurally Rational Players”, *American Economic Review* 88, 834-849.
- [13] Prendergast, C. (1999), “The Provision of Incentives in Firms”, *Journal of Economic Literature* 37, 7-63.
- [14] Rothschild, M. and J. Stiglitz (1970). “Increasing Risk: I. A Definition”. *Journal of Economic Theory* 2, 225-243.

- [15] 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.
- [16] Sloman, S. (2005). *Causal Models; How People Think about the World and its Alternatives*. Oxford University Press. New York.
- [17] Spiegel, R. (2004). “Simplicity of Beliefs and Delay Tactics in a Concession Game”. *Games and Economic Behavior* 47, 200-220.
- [18] Spiegel, R. (2006). “The Market for Quacks”. *Review of Economic Studies* 73, 1113-1131.
- [19] Spiegel, R. (2010). “Placebo Reforms”. Sapir Center Discussion Paper 3-11.
- [20] Stokey, N., R. Lucas and E. Prescott (1989), *Recursive Methods in Economic Dynamics*. Harvard University Press.

Appendix A: Permanent Shocks

Let us extend the model of Section 2 by assuming that every action a is characterized by an additional parameter $\rho_a \in [0, 1]$, and that the process (1) is modified as follows:

$$x^t = \begin{cases} x^{t-1} + \mu_{bt} + \varepsilon^t & \text{with probability } \rho_{bt} \\ x^{t-1} + \mu_{bt} + \varepsilon^t - \varepsilon^{t-1} & \text{with probability } 1 - \rho_{bt} \end{cases}$$

where ε^s is an independent draw from f_{b^s} . Thus, ρ_{bt} is the probability that the shock in period $t - 1$ is permanent; and just like the trend slope and the noise distribution, this probability is determined by the latest intervention prior to t .

I continue to assume that PMs have perfect information of the history. Thus, when player t moves, he is informed of the history $(x^0, a^0, \dots, x^{t-1}, a^{t-1}, x^t)$, as well as whether each of the shocks in periods $s \leq t$ were permanent. This means that the player knows ε^t and whether this shock is permanent. I analyze Markov perfect equilibrium with respect to the state (ε^t, z^t) , where $z^t \in \{0, 1\}$ and $z^t = 1$ means that ε^t is a permanent shock. I focus on Markov equilibrium not only because it is simpler and conventional in the literature, but also because it enables a perspective into the model that escapes the more general analysis of subgame perfect equilibrium.

When player t moves after a permanent shock, his payoff function is independent of the value of ε^t . Moreover, it is exactly as if he moved after a *transient* noise realization

$\varepsilon^t = 0$. Therefore, in Markov equilibrium, his behavior should be the same. As we will see, a player who moves after a transient shock follows a cutoff rule similar to that of the basic model. The exact expression for the cutoff will depend on whether it is positive (in which case future PMs would intervene after a permanent shock) or negative (in which case they would not).

Define the noise realizations ε^* and ε^{**} as the (unique) solutions of the following equations:

$$\begin{aligned} (1 - \delta)\varepsilon^* &= \max_{a \neq 0} \{ \mu_a - \delta[\rho_a \varepsilon^* + (1 - \rho_a)R_a(\varepsilon^*)] \} \\ (1 - \delta)\varepsilon^{**} &= \max_{a \neq 0} \{ \mu_a - \delta(1 - \rho_a)R_a(\varepsilon^{**}) \} \end{aligned}$$

Note that if $\varepsilon^* < 0$, then $\varepsilon^{**} < \varepsilon^*$.

Proposition 2 *Markov perfect equilibrium is characterized as follows:*

(i) *If $\varepsilon^* > 0$, then each player t chooses 0 whenever $\varepsilon^t \geq \varepsilon^*$ and $z^t = 0$, and an action in $\arg \max_{a \neq 0} \{ \mu_a - \delta[\rho_a \varepsilon^* + (1 - \rho_a)R_a(\varepsilon^*)] \}$ otherwise.*

(ii) *If $\varepsilon^* \leq 0$, then each player t chooses 0 whenever $\varepsilon^t \geq \varepsilon^{**}$ or $z^t = 1$, and an action in $\arg \max_{a \neq 0} \{ \mu_a - \delta(1 - \rho_a)R_a(\varepsilon^{**}) \}$ otherwise.*

Proof. Fix a Markov equilibrium, and let $V(\varepsilon^t, 0)$ be player t 's expected discounted equilibrium payoff at the state $(\varepsilon^t, 0)$, conditional on choosing optimally among all actions $a \neq 0$. Recall that if $V(\varepsilon^{t+1}, 0) > 0$, then player $t + 1$ intervenes at $(\varepsilon^{t+1}, 0)$, and therefore player t gets no credit on changes in x beyond period $t + 1$. On the other hand, if $V(\varepsilon^{t+1}, 0) \leq 0$, player t 's payoff conditional on ε^{t+1} being a transient shock is by definition weakly below $x^{t+1} - x^t + \delta V(\varepsilon^{t+1}, 0)$. Then, the following inequality holds for every ε^t :

$$V(\varepsilon^t, 0) \leq \sup_{a \neq 0} \left\{ \mu_a - \varepsilon^t + \delta \left[\rho_a \min(0, V(0, 0)) + (1 - \rho_a) \int \min(0, V(\varepsilon, 0)) f_a(\varepsilon) d\varepsilon \right] \right\} \quad (7)$$

The expression for the integrand exploits the zero-mean property of f_a . A priori, inequality (7) need not be binding, because an action that maximizes the R.H.S is not necessarily an action that player $t + 1$ would choose at a state $(\varepsilon^{t+1}, 0)$ conditional on playing $a \neq 0$. However, note that the solution to the maximization problem on the R.H.S is independent of ε^t . Therefore, any action that maximizes this problem is also an action that player $t + 1$ will find optimal at $(\varepsilon^{t+1}, 0)$ conditional on playing $a \neq 0$. Therefore, the inequality is binding and yields a recursive functional equation. It is

straightforward to verify that this equation meets Blackwell's sufficient condition for a contraction and therefore has a unique solution.

Let us guess a solution and show its consistency. It is clear from (7) that $V(\varepsilon, 0)$ is decreasing in ε , hence there is a noise realization $\bar{\varepsilon}$ for which $V(\bar{\varepsilon}, 0) = 0$, such that the player is indifferent between $a = 0$ and some optimal action conditional on $a \neq 0$, denoted a^* . Denote $\rho \equiv \rho_{a^*}$, $\mu \equiv \mu_{a^*}$, $k \equiv k_{a^*}$, $f \equiv f_{a^*}$.

We need to distinguish between two cases. First, assume that $\bar{\varepsilon} > 0$, hence $\min(0, V(0, 0)) = 0$. Guess that for every $\varepsilon \geq \bar{\varepsilon}$, $V(\varepsilon, 0) + \varepsilon$ is equal to some constant w . Since $V(\bar{\varepsilon}, 0) = 0$, $w = \bar{\varepsilon}$. Then, we can rewrite (7) at $\varepsilon = \bar{\varepsilon}$ as follows, exploiting the fact that the expectation of ε is zero, as well as the observation that $V(\varepsilon, 0)$ is decreasing in ε :

$$\begin{aligned} \bar{\varepsilon} &= \mu + \delta(1 - \rho) \int_{\bar{\varepsilon}}^k [V(\varepsilon, 0) + \varepsilon - \varepsilon] f(\varepsilon) d\varepsilon \\ &= \mu + \delta(1 - \rho) \int_{\bar{\varepsilon}}^k [\bar{\varepsilon} - \varepsilon] f(\varepsilon) d\varepsilon \\ &= \mu + \delta(1 - \rho) \left[(1 - F(\bar{\varepsilon}))\bar{\varepsilon} - \int_{\bar{\varepsilon}}^k \varepsilon f(\varepsilon) d\varepsilon \right] \\ &= \mu + \delta(1 - \rho) \left[(1 - F(\bar{\varepsilon}))\bar{\varepsilon} + \int_{-k}^{\bar{\varepsilon}} \varepsilon f(\varepsilon) d\varepsilon \right] \end{aligned}$$

Rearranging and plugging in (5), we obtain:

$$(1 - \delta)\bar{\varepsilon} = \mu - \delta[\rho \cdot \bar{\varepsilon} + (1 - \rho)R(\bar{\varepsilon})]$$

hence, $\bar{\varepsilon} = \varepsilon^*$. Since by assumption $\varepsilon^* > 0$, $V(0, 0) > 0$, hence the guess that players intervene after a permanent shock is consistent. To obtain the actions that players choose conditional on intervening, plug the characterization of $\bar{\varepsilon}$ we have obtained into (7). We have thus proved part (i) of the proposition.

The second case is where $\varepsilon^* \leq 0$, such that $\min(0, V(0, 0)) = V(0, 0)$. Using the same reasoning as in the first case, we obtain part (ii) of the proposition. In particular, since $\varepsilon^{**} \leq \varepsilon^*$, players choose $a = 0$ at the state $(0, 0)$, hence the guess that players choose $a = 0$ after a permanent shock is consistent with the constructed equilibrium. ■

It follows that in equilibrium, all PMs display a preference for reform strategies that are associated with permanent shocks. In particular, if two interventions have the same μ and f , PMs will prefer the one with the higher ρ . Moreover, it is evident from the PMs' equilibrium choices conditional on intervening that PMs become less

risk averse when they choose actions with a higher ρ . The intuition is that the adverse selection effect that lies behind the equilibrium risk aversion arises from the mean-reverting aspect of the process governing x^t , and therefore diminishes when this aspect is attenuated.

Note that the precise way in which the probability of permanent shocks affects a PM's trade-off between risk and return depends on whether he expects his successors to intervene after a transient null shock - or, equivalently, after a permanent shock. This in turn affects the precise cutoff noise realization that determines the timing of interventions. However, the qualitative effect of permanent shocks, as described in the previous paragraph, is independent of this distinction.

Appendix B: Pure Placebo Reforms under General Processes

Let us examine a specification of the model, in which x evolves according to an arbitrary finite-state Markov process, and PMs are unable to affect its course. Therefore, any intervention is by definition a Placebo reform. This is a stochastic generalization of the example given in the Introduction. Subgame perfect equilibrium in this case has a simple structure that illuminates the strategic considerations in our model.

Formally, assume that for every $t > 0$, $x^t = x^{t-1} + d^t$, where the initial condition is $x^0 = 0$, and d^t is governed by the following stochastic process. Let Q be a finite set of states, where $q^t \in Q$ denotes the state of the process at time t . The set of feasible actions is $A = \{0, 1\}$ for all players. Let $d(q)$ represent 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 $\tau(q^{t+1} | q^t)$ be the probability that the process switches to q^{t+1} in period $t + 1$ conditional on being in q^t in period t .

Proposition 3 *In subgame perfect equilibrium, each player t chooses $a = 1$ and earns a payoff of $V(q^t)$ if and only if $V(q^t) > 0$, where V is uniquely determined by the following recursive equation:*

$$V(q) = \sum_{q'} [d(q') + \delta \cdot \min(0, V(q'))] \cdot \tau(q' | q) \quad (8)$$

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 player t moves as a sequence of actions and states $h = (q^0, a^0, q^1, a^1, \dots, a^{t-1}, q^t)$. Let $q(h)$ be the state of the process at the history h , and

let $V^*(h | \sigma)$ be the expected payoff that player t attains if he chooses $a = 1$. Note that his equilibrium payoff is by definition

$$U(h | \sigma) = \max(0, V^*(h | \sigma)) \quad (9)$$

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:

$$V^*(h | \sigma) = \sum_{q'} [d(q') + \delta \cdot \min(0, V^*((h, 1, q') | \sigma))] \cdot \tau(q' | q(h)) \quad (10)$$

This equation is a contraction, and therefore has a unique solution. Furthermore, because periodic payoffs are only a function of q , the solution is measurable with respect to Q , and thus given by (8). Combining it with (9), we complete the proof. ■

The recursive function defined by (8) captures the essence of the PMs' strategic considerations in a pure placebo-reform setting. 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 .

The inverted search analogy, presented in Section 3 in the context of stationary growth processes, is made transparent by Equation (8). From player t 's perspective, future players behave as if they collectively solve a Markovian minimization problem. However, his own action selects the *maximum*, rather than minimum, between 0 and $V(q^t)$.

This equilibrium characterization is applicable to situations in which one PM confronts an irreversible reform decision, anticipating the future placebo reforms motivated by his successors' career concerns. Formally, suppose that player 0 could irreversibly - and invisibly, as far as the public is concerned - select a stochastic process (Q, q^0, d, τ) just before he chooses an action $a \in \{0, 1\}$. Then, this PM would select the process that generates the highest $V(q^0)$ in equilibrium (as long as this value is positive), rather than the process that maximizes the expected discounted value of x .