

Reform Fatigue ^{*}

T. Renee Bowen[†] Jackie M.L. Chan[‡] Oeindrila Dube[§]
Nicolas Lambert[¶]

March 28, 2016
(Preliminary and incomplete)

Abstract

We present a rational theory of reform fatigue. At each instant a politician chooses to divide effort between reforms and the status quo. This choice is modeled as a two-armed bandit problem. Reforms are expected to yield a higher rate of output to the voter than the status quo, conditional on the politician being competent. We interpret competence as the administrative ability to ensure successful implementation of reforms. The politician’s competence is unknown ex-ante to both the politician and the voter. In addition the voter is unable to observe the politician’s effort on reform, but only observes aggregate output. In equilibrium the voter gives the politician term lengths that depend on the timing of success. The politician experiments with reforms at the beginning of his first term, but gradually decreases the rate of reforms in the absence of early success. We call this gradual reduction in experimentation *reform fatigue*. The theory thus predicts that reform fatigue follows a political cycle. We provide empirical evidence of reform fatigue cycles in financial policies among presidential countries.

^{*}We thank Alberto Alesina, Jeremy Bulow, Allan Drazen, Frederico Finan, Chad Jones, Navin Kartik, Nicolas Klein, Marina Halac, Eddie Lazear, Condi Rice, Ken Shotts, Bruno Strulovici, Nancy Zimmerman and participants at the Stanford University Political Economy conference, the SITE 2014 session on Dynamics of Collective Decision Making, the Political Economy of International Organizations conference 2015, The Warwick/Princeton Political Economy Conference in Venice 2016, for helpful comments and conversations. This paper subsumes the papers circulated under the titles “Reform Fatigue in Economic Policy” and “Reform Cycles and Populist Cycles”, by Bowen, Chan and Dube, and “Reform Fatigue” by Bowen and Lambert.

[†]Stanford University

[‡]Chinese University of Hong Kong

[§]New York University

[¶]Stanford University

“No good deed goes unpunished.” -*Clare Boothe Luce*

1 Introduction

Many governments embark on market reforms, for example, liberalization of international trade, financial markets, and reduced regulation. However, it is often the case that the pace of these reforms diminishes over time, leaving intended reforms incomplete. This was true of many Latin American countries in the 1980s and 1990s, Turkey in the 1990s, Poland in the mid 2000s and, more recently, Greece. This phenomenon has come to be known in the media as “reform fatigue”.¹

A common explanation for reform fatigue is that, in the absence of improvements in outcomes, reforming policymakers fear that reforms will lose voters’ support, and consequently the policymakers themselves will be removed from office (see, for example, Lora et al., 2003). Financial reforms, such as those recommended by the International Monetary Fund (IMF), are often a source of tension between voters and elected officials, and politicians must make trade-offs between the political cost of implementing the reform and the perceived benefit to voters. Consistent with the political explanation, financial reforms appear to exhibit a political cycle. Figure 1 gives the average pace of reforms within a politician’s term in office. A term averages three years in the sample of countries we examine. In the figure, year 0 is the election year. The figure shows that, on average, a new politician quickly begins pursuing reforms at the beginning of a term in office, but decreases the pace of reforms thereafter. This pattern is consistent with reform fatigue as commonly described.

Although the term reform fatigue has been used in modern times with particular reference to financial reforms initiated by the IMF, such trade-offs are not restricted to IMF reforms, nor are they new. As an example, in December 1887 Grover Cleveland attempted to reduce high protective tariffs and subsequently lost his bid for re-election. He was quoted as saying “What is the use of being elected or re-elected unless you stand for something?”² In the case of Grover Cleveland, the choice of a policy whose benefits were not directly observable to the voters, but perceived to be beneficial by the executive, was electorally costly.

The political explanation for reform fatigue presents the following puzzle: if voters are cognizant of the stochastic nature of reform output and believe that incumbents are earnestly pursuing beneficial reforms, why would rational voters punish reform efforts at the polls? In this paper we offer a resolution to this puzzle based on voters’ beliefs about the politician’s

¹See ‘Why reform fatigue has hit the East’, *Financial Times*, September 26, 2006, and ‘IMF warns reform fatigue holding Greece back’, *Wall Street Journal*, June 2014.

²See Freidel and Sidey (2006).

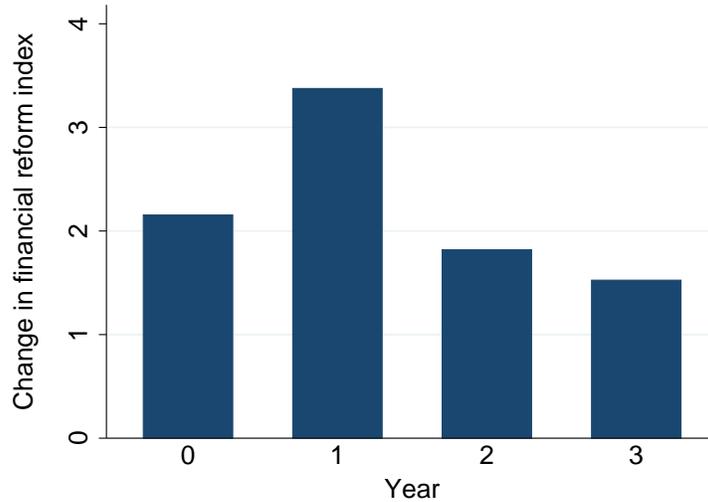


Figure 1: Average change in reforms within a term.

competence to conduct reforms. We posit that voters do not lose confidence in the reform *per se*, but lose confidence in the politician’s ability to successfully implement them. A politician may also be uncertain of his ability to conduct reforms, and thus experiments with reforms until he believes that he will be unable to successfully implement them. This paper thus presents a rational theory of reform fatigue.

The problem of reform fatigue is not restricted to the political context. Many organizations require the use of talented individuals who can produce success with relatively high frequency, but talent may be unknown *ex-ante* to both the organization and the individual. Much of the economics literature has focused on incentives to *select* the individual who knows he is talented. However, in many economic situations talent is proved only once the individual is on the job, and thus may not be known with certainty *ex-ante* to the individual or the organization. Examples include junior faculty at a university, professional athletes, entertainers, and mutual fund managers. In the case of junior faculty, the university would like to encourage junior faculty to take on ambitious projects, but the competence of the junior faculty member in executing such ambitious projects may be unknown to both the faculty member and the university. Many professional athletes enter their careers highly touted but do not “live up to the hype”. Yet it takes a coach some time to become sufficiently pessimistic about the athlete’s ability to dismiss the athlete. A difficulty common to these settings is identifying and rewarding uncertain talent that can be hard to distinguish from luck, when effort and the source of output is unobservable to the organization. While we find these applications interesting, the problem of retention (without the use of transfers) is most applicable in the

context of a politician seeking to engage in reforms, which is what we study.

We present a continuous time, infinite horizon model of reforms in the spirit of the career concerns literature and introduce experimentation.³ A politician is drawn at random and at each instant he can allocate a divisible unit of effort to the reform or the status quo. At any instant, the voter can choose to fire the politician at some cost, and is unable to commit to fire the politician at any given time. While in office, the politician obtains a flow payoff, but receives no payoff if fired, and is thus purely office-motivated. The voter gets a unit of benefit each time output is realized. Output may be generated from the reform or the status quo.

A politician entering office is either competent or incompetent. We regard competence as “administrative IQ”, as suggested by Rogoff (1990). In this respect, competence to conduct reform entails, for example, the ability to build coalitions to support and pass reform policy; write the text of legislation so that the enacted policy reflects the original intent; and ensure the successful implementation of the policy. A politician who is competent at reforming must be competent at each stage, and failure to execute may be interpreted as a failure of the reform. Consider the case of Peronist President Carlos Menem, who surprised everyone in 1991 by implementing a market liberalization program in Argentina. Program implementation required skilled political maneuvering. Menem had to gain the backing from powerful provincial governors using federal fiscal favors; and from key unions and business leaders with targeted privileges (Forteza et al., 2006). He also had to strategically invite key businessmen and politicians from the conservative party into his cabinet, to build a coalition that would support the passage of these reforms.

We take the perspective that effort exerted on reform does not deterministically translate into success. We thus model the politician’s choice to conduct reform or pursue the status quo as a two-armed bandit problem. The first arm is the “reform” (or risky) arm, and the second is the “status quo” (or safe) arm. Conditional on the politician being competent, the reform yields a unit of output (or, in the language of the bandit literature, a “success”) at a higher rate than the status quo per unit of effort. An example of reforms that were implemented on paper, but the actual impact was uncertain for some time were Mexico’s education reforms implemented in 2013. David Calderon, director of the education reform advocacy group Mexicans First commented on the reform saying, “Of course it’s just a change in the rules that still has to be turned into reality”.⁴ We assume that greater effort on reform translates into a higher probability of success, however effort by an incompetent politician

³Experimentation is modeled as in Keller et al. (2005). We describe our relationship to this literature in the literature review.

⁴See ‘Mexico Education Reform Passed By Senate, Looks To Remake Public School System’, *Huffington Post*, September 4, 2014.

will yield no success on the reform.

Consistent with the literature on career concerns, we assume that neither the politician nor the voter knows the politician's competence (or type) at the beginning of his term. Instead, the politician and the voter share a common prior belief about the politician's type. Thus, the politician, uncertain about his reform ability, must attempt reforms to learn his type. As an example, the Patient Protection and Affordable Care Act was signed into law in 2010 after passing several legislative hurdles, yet it met implementation challenges during the rollout of healthcare.gov, the website largely responsible for delivering insurance made possible by the act. It is reasonable to assume that a newly elected politician will be uncertain that he can be successful along every dimension, and so will the voter.

We assume that the allocation of effort to the reform and the status quo is unobserved to the voter. Much of the work to conduct reforms, for example organizing coalitions and writing legislation is largely unseen to ordinary citizens. Some of this effort can be observed to the interested researcher or reporter, but we argue that the cost of acquiring such information is prohibitive to the average person. Further, we assume that the voter will observe if a success occurs (for example voters observe an increase in gross domestic product), but whether it was the reform or the status quo that generated the success is unobserved to the voter. The arm that generated the success is observed only by the politician. When a success is observed by the voter it is not always clear if it was due to the reform or to luck while pursuing the status quo. For example, in the case of IMF suggested financial reforms, if personal incomes rise subsequent to the reform, it may be unclear to the voter if this was due to success of the financial reforms, or a positive income shock. An observed success on the reform means that the politician is competent for certain. Thus the voter and the politician may learn about the politician's competence gradually, but potentially at different rates. This reflects the reality that voters often re-elect a politician without being completely certain that he is competent, and a politician is likely to have more information about his competence than the voter by the end of a term in office.

To simplify the problem we consider an institutional setting in which the voter can commit to a success threshold and grants tenure to the politician if the success threshold is achieved.⁵ The voter however, is unable to commit to firing the politician at any given time, and at any moment can fire the politician before tenure is granted.⁶ Following much of the literature on strategic experimentation, the solution concept used throughout is Markov perfect equilibrium (MPE).

⁵One might also think of the tenure reward as the politician's legacy payoff.

⁶For comparison, we study the case of commitment in Section 5. The commitment case does not yield the gradual reform fatigue observed in practice, thus we think this case is less relevant to observed reforms.

In the simplest case of interest, where two successes are required, in equilibrium the voter gives the politician a fixed time to achieve the first success, and gives the politician an endogenous length of time (which can be thought of as a term) in which to generate the second success. This endogenous term decreases with the length of time it takes to achieve the politician's first success.

To see why the endogenous term length occurs in equilibrium, note that because of the voter's lack of commitment, the voter cannot learn the politician's type with certainty when there is a conflict of interest. A conflict of interest occurs when the belief about the politician's type is sufficiently low that the myopic best action is to choose the status quo, but the voter prefers further experimentation with reform because of her dynamic incentives. If the politician exerts no effort on the status quo in equilibrium, and is revealed to be good with certainty, then the voter will want to grant him tenure immediately. As a consequence, the politician will seek to maximize the immediate chance of success, and so will have an incentive to deviate to the status quo as this is myopically best. For the voter to remain uncertain, the politician must be induced to take an intermediate action, i.e., allocate resources to both the reform and the status quo at every instant. The voter must therefore keep the politician indifferent between the reform and the status quo at every instant when there is a conflict of interest. Since the politician can only be rewarded with more time in office, the voter uses the endogenous term length to provide sufficient flexibility to allow this indifference. Flexibility is required because beliefs, and hence payoffs, are evolving through time.

In equilibrium, a gradual decrease in the level of experimenting with reform is chosen if no success is observed for some period of time. This reflects the observed reform fatigue. If a success occurs during the period of intermediate experimentation, the voter is uncertain if the success was from the reform or the status quo, and the voter's belief about the politician's competence diverges from that of the politician. In this case, the voter may re-elect an incompetent politician with some probability, or conversely, may dismiss a competent politician. These will depend on the realization of the stochastic output.

As is true in many bandit problems with strategic choice, the politician under-experiments with reforms in equilibrium. By conducting the voter's optimal level of reform (a seemingly good deed), the politician is punished by the voter if no success is obtained. (As the opening quote by Clare Booth Luce says, 'no good deed goes unpunished'.) We show that reform fatigue is mitigated (i.e. reform effort increases) with a lower cost of firing, and the endogenous term length decreases. These results are intuitive, as with a lower cost of firing, the voter has a greater incentive to fire the politician, and thus the equilibrium induces higher effort from the politician.

We empirically investigate reform fatigue in financial policies, and show that the model

delivers the sort of rich dynamics we see in the data. We show that the cycle is present among countries with presidential systems. Countries with parliamentary systems, where executives are not elected directly, exhibit no such cycle. This is consistent with our theory, which presumes politicians are directly accountable to voters. The theory predicts that, conditional on being close to the end of the term, reforms may increase in response to an increase in output. We also see that this is true in the data. As the theory does not rely on an external organization to generate reform fatigue, we investigate the role of the IMF and demonstrate that the reform fatigue cycle is not significantly larger when countries participate in an IMF program. Reform fatigue appears to be a broader phenomena, present among both program participants and non-participants.

Literature review

There are at least two competing explanations for reform fatigue. One is that the benefits of the reform to various constituencies are uncertain and potentially uneven. When information about reforms are revealed and a sufficiently large constituency expects to lose from reforms, they will oppose those reforms. This explanation has been studied by Fernandez and Rodrik (1991) and more recently by Strulovici (2010). Another explanation is that there are different types of reforms with varying degrees of difficulty and reformers enact “easy” reforms in the beginning and are simply unable to enact more difficult reforms later on, hence reforms appear to cease. This *gradualism* in reforms has been explored by Dewatripont and Roland (1992) and Dewatripont and Roland (1995). Unlike Fernandez and Rodrik (1991), Dewatripont and Roland (1995) not only consider the ex-ante choices, but the choices of the median voter after the realization of the outcome from initial reforms as the median voter learns about the reform. These explanations are appealing, but we show empirically that reform fatigue follows a political cycle, a prediction absent in these theories.⁷

This paper is related to the substantial body of political economy research studying political failures, which was first articulated by Besley and Coate (1998). In this paper we show how a lack of commitment from the voter, combined with asymmetric learning, can undermine incentives for efficient experimentation on reforms by a politician. A similar question is explored theoretically in Canes-Wrone et al. (2001) and empirically in Canes-Wrone and Shotts (2004) within the context of pandering. Our contribution is to study the evolution of reform policies throughout a politician’s term in office, as he learns about his own competence. We show that the politician’s incentive to learn about his type induces more effort on reform early in a term, but in the absence of output leads to a decrease in

⁷Tornell (1998) also provides a theory of reform, but does not focus on the electoral timing of reform.

effort on reform. The decrease in reform effort is driven by the voter’s inability to commit, and hence the need for the politician to obscure learning about his type when there is a conflict of interest.

We study a politician concerned about retaining his position, thus, our paper is related to the large career concerns literature, which is recently summarized by Ashworth (2012).⁸ Our theory is closest to Jackson and Aghion (Forthcoming) who also consider the problem of motivating a politician through firing incentives, and when there is learning about the quality of the politician. There are two main differences in Jackson and Aghion (Forthcoming). First we include the problem of hidden actions and hidden types, so the choice of the politician and the politician’s true beliefs are unobserved to the voter. Second, in our model, the competence of the politician is related to his ability to deliver on reforms, rather than his ability to perfectly observe the random state of the world.

There is a large literature on bandit problems in economics including the classic work of Keller et al. (2005), however, few papers have incorporated hidden actions, and asymmetric learning.⁹ One notable exception is Halac et al. (Forthcoming), which has several differences with the current paper. Halac et al. (Forthcoming) are interested in an optimal monetary contract, whereas we are interested in a setting where the voter’s only means of creating incentives is to retain or replace the politician. In other words, we consider that wages are fixed, and the “contract” that the voter can offer is a firing contract—a somewhat blunt tool.¹⁰ Second, in Halac et al. (Forthcoming) the politician knows his type prior to beginning the project, hence learning is only about the quality of the project.¹¹ A small number of authors have applied the tools of the bandit literature to the study of reforms, including Strulovici (2010). Similar to our work, Strulovici (2010) considered reforms as risky experiments. As mentioned, unlike our work, Strulovici (2010) considers that reforms have heterogeneous effects on voters that are learned over time, and the theory does not predict a reform cycle.¹²

Empirical studies have also advanced reasons for reform to be higher when new governments

⁸This literature includes Alesina and Cukierman (1990), Besley and Case (1995), Banks (1998), Holmstrom (1999), Ashworth (2005), Berliant and Duggan (2010), DeMarzo and Sannikov (2014), Bonatti and Hörner (Forthcoming), Duggan and Forand (2014), Besley and Case (1995) and Canes-Wrone et al. (2001). These papers do not consider political accountability for reforms, with reforms modeled as a bandit problem.

⁹Recent contributions to the experimentation literature include Garfagnini (2011), Hörner and Samuelson (2013) and Guo (Forthcoming).

¹⁰See Bowen and Mo (Forthcoming).

¹¹In this sense, it is closer in spirit to Klein (Forthcoming). However in Klein (Forthcoming) there is a single politician, and the voter’s objective is to induce the politician to always experiment. In our setting always experimenting is not always optimal for the voter, in some cases the voter would prefer if the politician stopped experimenting.

¹²Other contributions to the literature on reforms include Lizzeri and Persico (2011), Callander and Harstad (2015). These do not model reforms as a bandit problem.

first come to office. For example, some have suggested that new incumbents will want to undertake reforms right away so that they can realize gains prior to the next election, assuming that gains may be delayed. This is also known as the “honeymoon hypothesis” (Krueger, 1993; Abiad and Mody, 2005). Abiad and Mody (2005) investigate if the tendency to reform is higher during the incumbent’s first year in office, as a part of their broader exploration of what determines financial reforms. They find partial support for this hypothesis, in that reforms are typically higher during the first year of the electoral term, although this effect is imprecisely estimated. Consistent with Abiad and Mody (2005) we find a significant positive effect on reform during the first year after the election, suggesting more reforms are implemented in the beginning of the electoral term. As we are concerned with investigating the entire dynamic path of the elected official’s time in office, we also demonstrate a significant negative effect in the year prior to the next election, suggesting that the tendency to reform is diminished at the end of the electoral term.¹³ This pattern is consistent with the rational theory of reform fatigue we present. We are unaware of any papers that theoretically examine the honeymoon hypothesis with rational voters.

A number of papers have empirically examined the relationship between reforms and other outcomes, such as growth, the level of democracy, or labor market performance (e.g., Christiansen et al., 2013; Giuliano et al., 2013; Di Tella and MacCulloch, 2005; Feldmann, 2012). However, none of these studies focus specifically on political cycles. Lora et al. (2003) also use public opinion questions from to examine support for market reforms. They define reform fatigue as falling public support for reforms, and document this type of reform fatigue, over time, in Latin America. Trend decreases in public support for reforms is not inconsistent with our account. However, our focus lies in examining how the tendency to reform varies over the course of the electoral cycle.

There are a significant number of papers studying political cycles, including the seminal works of Nordhaus (1975) and Rogoff (1990). The political budget cycle has been summarized and well documented in Drazen (2001) as well as Brender and Drazen (2005) and Shi and Svensson (2006); and, a political aid cycle has been documented in Faye and Niehaus (2012). More recent work in this literature includes Canes-Wrone and Shotts (2004) and Ales et al. (2014).

¹³Whereas Abiad and Mody (2005) cover both parliamentary and presidential regimes between 1973-1996, and have 805 observations, we focus on only presidential regimes and have 1195 observations covering 1975-2005. Another difference with Abiad and Mody (2005) is that we use data from Abiad et al. (2008). The data used in Abiad and Mody (2005) has six rather than seven dimensions. It excludes securities market policy and prudential regulations, but includes a measure of operational restrictions. These differences are discussed in detail in Abiad et al. (2008). Our cycle is more precisely estimated than the first year effect in Abiad and Mody (2005).

The remainder of the paper is organized as follows. In Section 2 we present our stylized model of reforms with experimentation and private information. Section 3 discusses the first best level of experimentation for the voter. Section 4 constructs the MPE with reform fatigue. Section 5 gives the MPE for the case in which the voter can commit to a firing policy. Section 6 provides empirical evidence of reform fatigue and Section 7 concludes.

2 Model

We present a stylized model of a politician choosing reforms versus the status quo under the shadow of electoral incentives. A voter (she) and an incumbent politician (he) interact during the politician’s time in office. Time is continuous and the horizon is infinite.¹⁴ The voter and politician discount the future at a common rate r . The politician has a type, which is either competent or incompetent. We denote by $\theta \in \{0, 1\}$ the politician’s type, where the politician is competent if $\theta = 1$. The politician and the voter share the common prior belief that the politician is competent with probability q_0 .

At every instant that he is in office, the politician must choose to divide one unit of work resource between two tasks—reform and the status quo.¹⁵ The tasks are modeled as a two-armed exponential bandit as in the classic work of Keller et al. (2005). If the politician works x units on the reform during a small period $[t, t + dt)$ (and so works $1 - x$ on the status quo during that same period), the reforms generate one unit of output with probability $1 - e^{-\lambda_r \theta x dt} \approx \lambda_r \theta x dt$, and the status quo generates one unit of output with probability $1 - e^{-\lambda_s (1-x) dt} \approx \lambda_s (1-x) dt$. If a unit of output is generated on either reform or status quo, we say that a success has occurred. The probability of successes are independent across time and tasks. We assume $\lambda_r > \lambda_s$ so that the reform delivers output at a higher rate than the status quo, and $q_0 > \lambda_s / \lambda_r$ so that experimentation is myopically optimal for the voter at the beginning of the term.

¹⁴We seek to explain a gradual reduction of effort on reform, and for this reason the equilibrium we characterize requires continuous time. For the politician to exert an intermediate level of effort on reform, due to the voter’s lack of commitment, the politician must be indifferent between the reform and the status quo for the interval of time in which the gradual reduction occurs. If time is discrete, then there is not sufficient flexibility to induce this indifference. The politician can be made indifferent for at most one period for a fixed set of parameters with discrete time. A similar problem to ours is studied by Manso (2011) with two discrete periods. Unlike our setting, Manso (2011) assumes that the agent has a cost of effort, the principal has flexibility to offer performance-contingent payments, and, importantly, the principal is able to commit. In a two-period model, without commitment and without the flexibility of giving the politician a bonus, the only equilibrium when there is a conflict of interest is when the politician exerts effort only on the status quo.

¹⁵Equivalently, in the interpretation of Hörner and Samuelson (2013), the politician randomizes over the status quo and reform.

At any instant the voter can decide to fire the politician. We assume that, in general, the voter cannot commit to fire or retain the politician at any point in time, thus the voter makes her choice based on the information she has available at the time.¹⁶ When she fires the politician, the voter gets a lump sum payment λ_p/r . Alternatively, we might consider that when the voter fires the current politician she has to pay a cost C and gets a new politician with prior q_0 . For every $C > 0$, there exists some λ_p such that the lump sum λ_p/r corresponds to the continuation payoff of getting a new politician while incurring firing cost C . We assume $\lambda_s < \lambda_p < \lambda_r$. In addition we assume that the voter can commit, up front, to a number of successes N after which the politician cannot be fired—i.e., the politician is granted tenure. If $N = \infty$, then the voter never gives tenure to the politician.

The voter values successes independently of how they are generated. That is, the voter values output whether it comes as a result of reform, or good luck with the status quo. Each success gives the voter a payoff normalized to 1.¹⁷ She gets zero payoff the rest of the time the politician is hired. The politician gets a flow payoff of 1 per unit of time during the time he is hired. The politician, thus, only cares about retaining office.

The politician observes the successes as they occur and observes which task generates them. The politician's actions are hidden, that is, the voter does not observe how the politician divides his work between the status quo and reform. The voter observes successes, but does not know where they came from. Neither the politician nor the voter observes the politician's type ex-ante, rather, they learn about it over time.

3 First best

We present the best solution for the voter if all information is observable, and the voter can dictate the politician's action. In the solution to the first best, the voter would prefer to fire the politician and obtain the continuation payoff λ_p/r rather than switch to the status quo and obtain the continuation payoff λ_s/r . Thus the voter's optimal solution is equivalent to the first best solution with a two-armed bandit, for which the safe arm generates output with probability λ_p . As in Keller et al. (2005), the voter sets a stopping time T^* to fire the politician who has not gotten any success. At any time the politician is hired, the politician will put full effort on reform.

The time T^* is found by making the voter indifferent between firing the politician at T^* or keeping the politician one more instant. Let q_t be the politician's belief at $t \leq T^*$ if the

¹⁶We consider the case in which the voter can commit in Section 5.

¹⁷It is also possible that payoff amounts are generated randomly, but we normalize payoffs to 1 for simplicity.

politician has not gotten any success. This is given by

$$q_t = \frac{q_0 e^{-\lambda_r t}}{q_0 e^{-\lambda_r t} + 1 - q_0}.$$

The voter's indifference conditions is, to the first order in dt ,

$$\frac{\lambda_p}{r} = \lambda_r q_{T^*} dt \left[1 + \frac{\lambda_r}{r} \right] + (1 - \lambda_r q_{T^*} dt - r dt) \frac{\lambda_p}{r}.$$

This indifference condition characterizes the first best stopping time T^* and we denote the first best threshold belief as $q^* = q_{T^*}$. We summarize the first best in the following proposition.

Proposition 1. *In the first best solution, the politician puts full effort on reform until the belief that the politician is good reaches the threshold belief*

$$q^* = \frac{\lambda_p}{\lambda_r + \frac{\lambda_r}{r}(\lambda_r - \lambda_p)}.$$

If no success is obtained before the belief reaches q^ , then the politician is fired, but if a success is obtained before q^* , then the politician is retained forever and puts full effort on reform thereafter.*

The proof of Proposition 1 follows from Keller et al. (2005).

We note that the first best can be obtained in an equilibrium if the voter can observe the politician's action, even if the source of output is unobservable. The voter could simply fire the politician before T^* if the politician deviates from full effort on reform. Similarly, if the source of output is observable and effort is unobservable, then the voter can achieve the first best by giving tenure to the politician if a success is observed on the reform arm before T^* , and firing the politician at T^* otherwise. Thus with observable actions or observable source of output the first best is achievable and no reform fatigue is observed. A politician who does not choose reform, or does not have success on reform sufficiently early is fired.

There can be no conflict of interest between the voter and politician if the myopic threshold belief for the politician λ_s/λ_r is lower than the optimal threshold belief for the voter q^* . Hence, we assume from now on that $\lambda_s/\lambda_r > q^*$, which is the necessary and sufficient condition for a conflict of interest to occur.

4 Markov perfect equilibrium

We focus on Markov perfect equilibria, following much of the literature on strategic experimentation. The state variables are, for the voter, the probability that the politician is competent, and the distribution over the politician's belief. For the politician, the state variables are the voter's state variables, and the probability that he is competent given his own information.

An equilibrium is a reform policy $\chi_t : [0, 1] \times \mathbb{B} \rightarrow [0, 1]$ for the politician and a firing policy $\Upsilon_t : \mathbb{B} \rightarrow \{0, 1\}$ for the voter, where \mathbb{B} is the space of beliefs of the voter. Beliefs are updated via Bayes' rule. Note that the politician may have more information than the voter and, as a result, the beliefs of the politician and the voter about the politician's competence may diverge.

Consider the simplest case of $N = 1$, that is, the institution is such that the voter is required to tenure the politician after a single success. Since the prior belief is above the myopic threshold $q_0 > \lambda_s/\lambda_r$ we can show that a continuum of equilibria exist in which the politician pursues reform starting at time zero and for any length of time up to when the politician's belief q_t is at the myopic threshold, and then switches to the status quo until the first success. The voter will keep the politician after the first success, regardless of where the success came from per the institutional requirement. After the first success, the politician will pursue the voter's first best. It is clear that the voter will have an incentive to tenure the politician when the first success occurs since at that time $q_t \geq \lambda_s/\lambda_r > q^*$. Furthermore, the voter will have no incentive to fire the politician before the belief is at the myopic threshold, because $\lambda_s/\lambda_r > q^*$. The voter thus has no incentive to deviate. The politician has no incentive to deviate, since any deviation is unobserved to the voter, hence cannot affect his probability of re-election. In these equilibria (and any equilibrium when $N = 1$) the politician never exerts effort on both the reform and the status quo at the same instant, and thus will have no information that is private regarding his type. We are interested in equilibria in which the politician may be fired, and is able to use his private information about his type. We explore the simplest version of this next.

4.1 Equilibrium with reform fatigue

We seek an equilibrium in which the voter fires the politician with some probability on the equilibrium path and in which the politician uses his private information about his type. We ask if such an equilibrium exhibits reform fatigue, in the sense of a gradual reduction of effort on reforms. Since the case with one success does not exhibit a gradual reduction, we consider the next simplest case in which the voter commits to giving the politician tenure

after 2 successes. We have the following result.

Proposition 2. *For $N = 2$, there exists a Markov perfect equilibrium in which the politician is fired on the equilibrium path with some probability. This equilibrium exhibits reform fatigue. That is, effort on reform is gradually reduced if there is no early success.*

The proof of the proposition follows from the construction of the equilibrium. There exists a continuum of equilibria with reform fatigue, but we focus on the one that maximizes the payoff of the voter.

We describe the equilibrium briefly and then provide a more precise construction below. In equilibrium, conditional on no success being observed, a politician puts full effort on reform until the myopic threshold belief is reached at time T_m . The politician is fired at a time $T_1 \geq T_m$ if no success has occurred before T_1 . The politician who got a first success at time $t \in (T_m, T_1)$ will get a length of time Δ_t to get a second success, and if he fails to do so, he is fired. If he succeeds, he is granted tenure. Denote τ as the time as which the first success occurs, then $T_2 = \tau + \Delta_\tau$ is the time at which the politician is fired if there is no second success. The equilibrium thus consists of three phases: phase I is before the first success, phase II is after the first success and before the second, and phase III is after the second success. These phases are illustrated in Figure 2. We proceed by backward induction, beginning with phase III.

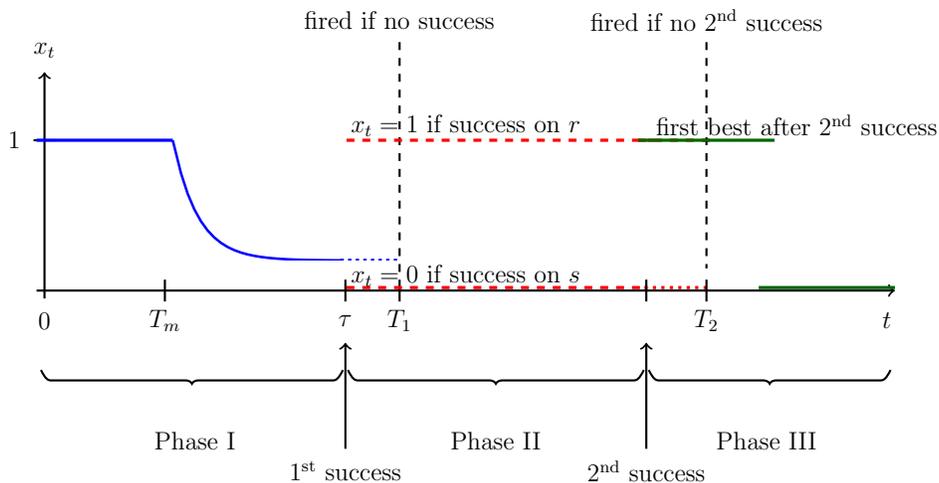


Figure 2: Equilibrium phases

Phase III: After the second success

After the second success the voter's criterion is met, the politician is granted tenure and there is no conflict of interest. In the equilibrium we are looking for, the politician does what is best for the voter when there is no conflict of interest. Let $H(q_t)$ be the voter's continuation value after the second success when the politician has belief q_t . Following Keller et al. (2005) in the first best there is a cutoff belief \bar{q} given by

$$\bar{q} = \frac{\lambda_s}{\lambda_r + \frac{\lambda_r}{r}(\lambda_r - \lambda_s)},$$

such that below the cut-off it is optimal to put full effort into reform and above the cutoff it is optimal to put full effort into the status quo. Note that we assume the politician is granted tenure after the second success, so the voter cannot fire the politician to obtain λ_p/r if the politician is sufficiently pessimistic about his type. The best thing for the voter is for the politician to switch to the status quo. The politician will do this even though effort is unobserved, because he is indifferent between all actions at this point. If the belief about the politician's type is q_t , then the value to the voter after a second success is thus

$$H(q_t) = \begin{cases} \frac{1}{r} \left[\lambda_r q_t + (\lambda_s - \lambda_r \bar{q}) \left(\frac{1-q_t}{1-\bar{q}} \right) \left(\frac{(1-q_t)\bar{q}}{(1-\bar{q})q_t} \right)^{\frac{r}{\lambda_r}} \right] & \text{if } q_t > \bar{q} \\ \frac{\lambda_s}{r} & \text{if } q_t \leq \bar{q}. \end{cases} \quad (1)$$

Note that in equilibrium the politician's belief will not fall below \bar{q} before the second success, and thus $H(q_t)$ is given by the first line of equation (1) in equilibrium.

Phase II: After first success, before second success

After the first success, but before the second success, the politician knows that he is retained forever if he gets a second success. The politician therefore has a myopic incentive to get a second success as quickly as possible, and there is no future benefit from experimentation. If the politician got the first success from the reform, the politician knows that he is competent and will put full effort on reform thereafter. If the politician got the first success on the status quo, the politician does not know if he is competent, and will put full effort on the status quo until the second success.¹⁸

Recall that Δ_t is the time the voter has given to the politician to get a second success conditional on the first success occurring at time t . If the politician was exerting effort on

¹⁸This assumes his belief is below his myopic threshold λ_s/λ_r which will be true in equilibrium.

the reform and the status quo at the time the first success was obtained, the voter's belief will diverge from the politician's belief. Let τ be the time of the first success and q_τ is the politician's belief that the politician is good at time τ . Note that given the conjectured equilibrium, q_τ is also the voter's belief that the politician is good at time τ , and is thus known to the voter. Let x_τ be the action played by the politician at that time, and let p_t be the voter's belief at time $\tau + t$ that the first success was obtained on the reform. Note that this is different from the voter's belief that the politician is good. By Bayes' rule we have

$$p_0 = \frac{\lambda_r x_\tau q_\tau}{\lambda_r x_\tau q_\tau + \lambda_s (1 - x_\tau)}. \quad (2)$$

Thus at the instant the success is obtained, the voter forms the belief p_0 that the success was obtained on the reform. By Bayes' rule, this belief is the probability of obtaining a success on reform as a fraction of the total probability of observing a success.

Recall that once the first success is obtained (but before the second success), the politician will put full effort on reform if the success was obtained from the reform, and otherwise will put full effort on the status quo in equilibrium. Thus at any time t after the first success, the voter is uncertain about the politician's actions. Conditional on observing no success, the voter updates the probability that the success was obtained on reform according to

$$p_t = \frac{e^{-\lambda_r t} p_0}{e^{-\lambda_r t} p_0 + e^{-\lambda_s t} (1 - p_0)}.$$

We use these beliefs and the voter's indifference condition at time $t + \Delta_t$ to derive the politician's effort at any time t . This is summarized in the next lemma.

Lemma 1. *If a first success occurs at time τ then the effort exerted on reform at time τ is*

$$x_\tau = \left[1 - e^{-(\lambda_r - \lambda_s)\Delta_\tau} \frac{\lambda_r q_\tau (\lambda_r - \lambda_p) (1 + \lambda_r/r)}{\lambda_s [\lambda_s (1 + H(q_\tau)) - \lambda_p (1 + \lambda_s/r)]} \right]^{-1}. \quad (3)$$

The proof of Lemma 1, and henceforth all proofs, are included in the Appendix unless otherwise indicated. From Lemma 1 we have the effort on reform before the first success x_τ as a function of beliefs q_τ and the evaluation period for the second success Δ_τ .

The next lemma gives the voter's expected continuation payoffs after the first success.

Lemma 2. *If the first success is obtained at time τ , then the voter's expected continuation payoff if the politician obtained the first success on reform is*

$$V_\tau^R = \frac{\lambda_r}{r} - \left[\frac{\lambda_r}{r} - \frac{\lambda_p}{r} \right] e^{-(r+\lambda_r)\Delta_\tau}, \quad (4)$$

and the voter's expected continuation payoff if the politician obtained the first success on the status quo is

$$V_\tau^S = \frac{1 + H(q_\tau)}{1 + r/\lambda_s} - \left[\frac{1 + H(q_\tau)}{1 + r/\lambda_s} - \frac{\lambda_p}{r} \right] e^{-(r+\lambda_s)\Delta_\tau}. \quad (5)$$

Phase I: Before the first success

Let q_t be the politician's belief that he is competent at time t , before the first success. At all times $t \leq T_m$ there is no conflict of interest, and the politician will put full effort on reform. If $t \leq T_m$ then $\Delta_t = \infty$. At all times $T_m < t \leq T_1$ there is a conflict of interest. The unsuccessful politician will play action $x_t \in (0, 1)$. For this intermediate action to be an equilibrium, the politician has to be indifferent between the status quo and reform. We first calculate the politician's continuation payoff W_t at any time t in phase I.

Lemma 3. *The politician's continuation payoff for a fixed Δ_t is*

$$W_t = e^{(\lambda_s+r)t} \int_{T_1}^t \frac{\lambda_s [e^{-(\lambda_s+r)(z+\Delta_t)} - e^{-(\lambda_s+r)z}]}{r} dz. \quad (6)$$

We calculate Δ_t next. Consider $t \in (T_m, T_1)$ when the politician is indifferent between the reform and status quo. If at time t the politician's belief is q_t and his continuation payoff is W_t , then the equilibrium Δ_t solves

$$\lambda_r q_t [W_t^R - W_t] = \lambda_s [W_t^S - W_t],$$

which is the politician's indifference condition. We can show that $W_t^S = \frac{1}{r} [1 - e^{-(\lambda_s+r)\Delta_t}]$ and $W_t^R = \frac{1}{r} [1 - e^{-(\lambda_r+r)\Delta_t}]$.¹⁹ Substituting W_t^R and W_t^S into the above expression and rearranging gives

$$\lambda_r q_t \left[\frac{1}{r} - \frac{1}{r} e^{-(\lambda_r+r)\Delta_t} \right] - \lambda_s \left[\frac{1}{r} - \frac{1}{r} e^{-(\lambda_s+r)\Delta_t} \right] = (\lambda_r q_t - \lambda_s) W_t. \quad (7)$$

This condition determines Δ_t given values for q_t and W_t .

To solve for the equilibrium values before T_1 , we first solve for the boundary conditions at T_1 . Since the politician is fired at T_1 we know that $W_{T_1} = 0$. This along with the voter's indifference condition, allows us to solve for the boundary values. These are given in the next lemma.

¹⁹The derivation of W_t^S is given in the proof of Lemma 3 in the Appendix. The derivation of W_t^R is analogous.

Lemma 4. *The values q_{T_1} , x_{T_1} , and Δ_{T_1} are implicitly given by the following three equations:*

$$q_{T_1} = \frac{\lambda_s [1 - e^{-(r+\lambda_s)\Delta_{T_1}}]}{\lambda_r [1 - e^{-(r+\lambda_r)\Delta_{T_1}}]}, \quad (8)$$

$$x_{T_1} = \frac{\lambda_p - \lambda_s(1 + V_{T_1}^S - \lambda_p/r)}{\lambda_r q_{T_1}(1 + V_{T_1}^R - \lambda_p/r) - \lambda_s(1 + V_{T_1}^S - \lambda_p/r)}, \quad (9)$$

and

$$\left[1 - \frac{\lambda_r(1 + \lambda_r/r) - \lambda_p(1 + \lambda_r/r)}{\lambda_s(1 + H(q_{T_1})) - \lambda_p(1 + \lambda_s/r)}\right]^{-1} = \left[1 + e^{(\lambda_r - \lambda_s)\Delta_{T_1}} \frac{\lambda_s(1 - x_{T_1})}{\lambda_r x_{T_1} q_{T_1}}\right]^{-1}. \quad (10)$$

We can now calculate equilibrium values for any value of t . We first note that the law of motion for beliefs is $q'_t + q_t(1 - q_t)\lambda_r x_t = 0$. Using this we can calculate all equilibrium values for $t \in (T_m, T_1)$. This is summarized in the next lemma.

Lemma 5. *For $t \in (T_m, T_1)$, the equilibrium values q_t , W_t solve the system of differential equations*

$$\begin{aligned} q'_t &= -q_t(1 - q_t)\lambda_r x_t, \\ W'_t &= (\lambda_s + r)W_t - \frac{\lambda_s}{r} [1 - e^{-(\lambda_s+r)\Delta_t}], \end{aligned} \quad (11)$$

where the equilibrium value $\Delta_t = \Delta(q_t, W_t)$ is defined as the solution to

$$\lambda_r q_t \left[\frac{1}{r} - \frac{1}{r} e^{-(\lambda_r+r)\Delta_t} \right] - \lambda_s \left[\frac{1}{r} - \frac{1}{r} e^{-(\lambda_s+r)\Delta_t} \right] = (\lambda_r q_t - \lambda_s)W_t, \quad (12)$$

and the equilibrium value x_t is defined as

$$x_t = x(\Delta_t, q_t) = \left[1 - e^{-(\lambda_r - \lambda_s)\Delta_t} \frac{\lambda_r q_t (\lambda_r - \lambda_p) (1 + \lambda_r/r)}{\lambda_s [\lambda_s(1 + H(q_t)) - \lambda_p(1 + \lambda_s/r)]} \right]^{-1}. \quad (13)$$

The boundary conditions are $W_{T_1} = 0$ and q_{T_1} as defined by (8).

In general, (12) does not admit a closed form solution, except for some special cases.²⁰

²⁰For example, if $\lambda_r + r = 2(\lambda_s + r)$, then solving (12) reduces to solving a quadratic equation. Eliminating the irrelevant root, we obtain

$$\Delta_t = \frac{1}{\lambda_s + r} \log \left(\frac{2\lambda_r q_t}{\lambda_s - \sqrt{4\lambda_r q_t (rW_t - 1)(\lambda_s - \lambda_r q_t) + \lambda_s^2}} \right).$$

Plugging the expressions of Δ_t and x_t back into (11), we obtain an explicit ordinary differential equation.

The proof of Lemma 5 follows from Lemmas 1-4 and equation (7). This completes the construction of the equilibrium.

4.2 Discussion

To aid the discussion of the equilibrium, we solve the system of equations numerically and obtain results as illustrated below in Figure 3. The parameter values used in Figure 3 are $\lambda_r = 0.3$, $\lambda_s = 0.1$, $\lambda_p = 0.15$ and $r = 1$. For these parameter values $T_m = 0.207$ and $T_1 = 1$. Thus if a success is obtained before T_m the voter is certain that the politician is good and will keep the politician. That is, $\Delta_t = \infty$ for all $t \leq T_m$, and if no success is observed before T_1 , then the politician is fired.

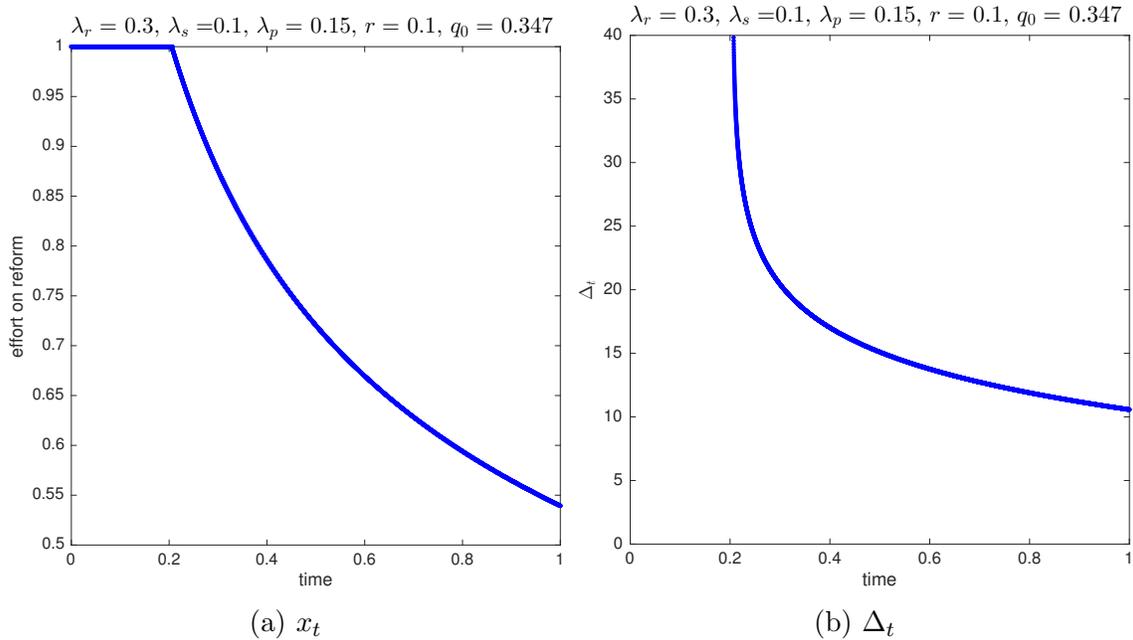


Figure 3: Equilibrium values assuming no success, $T_1 = 1$

If a success occurs in the interval (T_m, T_1) the voter is still uncertain about the politician's competence, but updates his belief positively. We provide an example of a success occurring in this interval in Figure 4. In Figure 4 we plot the voter's belief that the politician is good conditional on observing a success. Denote the voter's belief at time t that the politician is good by p_t^g . This is

$$p_t^g = p_t + (1 - p_t)q_r.$$

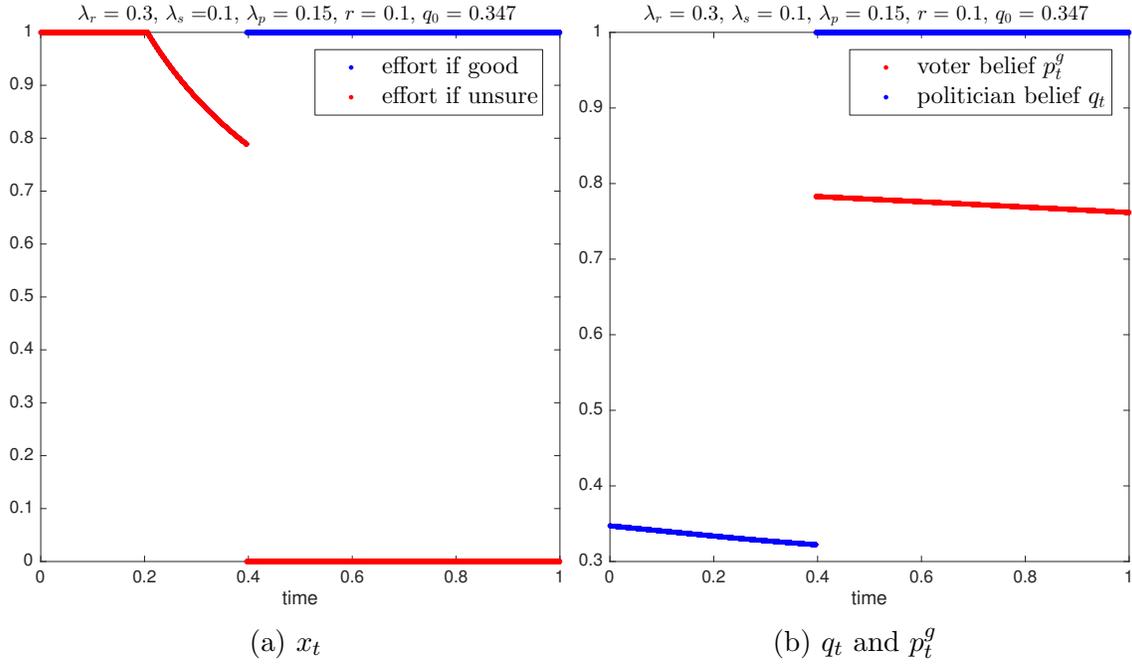


Figure 4: Equilibrium values assuming one success, $T_1 = 1$

In Figure 4, a success occurs at $\tau = 0.34$ and, assuming that the success was obtained from the reform, the politician’s belief jumps to 1, but the voter’s belief does not jump all the way to 1. (In the event that the success was obtained from the status quo, the politician’s belief would evolve in the same way it had prior to the success, and thus would continue to decrease and be lower than the voter’s belief.) Furthermore, the voter’s belief will decrease if there is no other success, and the voter is never certain that the politician is good.

Note that in the example of Figure 4 there is under-experimentation. That is, the first best belief at which the politician should be fired is $q^* = 0.20 < q_{T_1} = 0.32$. Thus, when there is no commitment, the voter fires the politician sooner than is optimal in equilibrium assuming no success. This is generally true in equilibria with reform fatigue.

An increase in the prior belief q_0 is illustrated in Figure 5. Since the equilibrium values at T_1 do not depend on q_0 , increasing q_0 simply increases the time it takes to reach the myopic threshold λ_s/λ_r and, correspondingly, q_{T_1} . Thus increasing q_0 increases the time before which the politician is fired if there is no first success. This is intuitive, as a politician believed to be competent with high probability, will be given more “wobble room” than one who is believed to be incompetent. This is summarized in the next lemma.

Lemma 6. *The time at which the voter fires the politician after no success T_1 is increasing in the prior belief that the politician is good q_0 .*

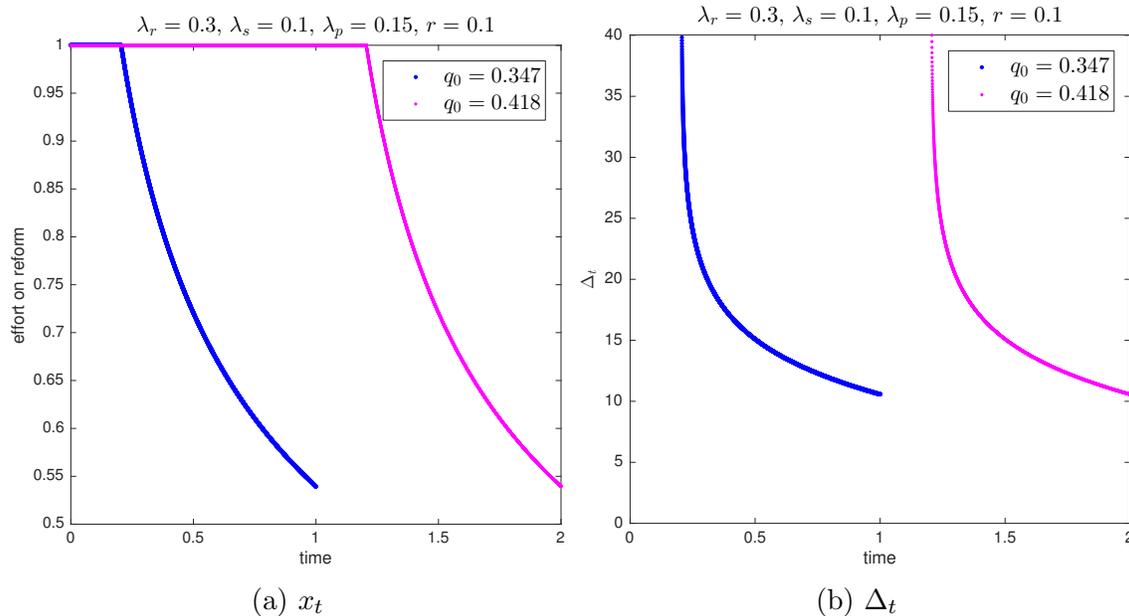


Figure 5: Increasing q_0

An increase in the rate of return from the reform λ_r is illustrated in Figure 6. Panel (a) shows that the politician experiments longer with the reform when the rate of return on the reform is higher. Panel (b) shows that the voter also gives the politician more time to implement the reform because the future value of a success on the reform is higher. Intuitively, the belief that the politician is good decreases at a faster rate with a higher value of λ_r , because more effort is being exerted on the reform, and the rate of updating is higher. We summarize this in Lemma 7.

Lemma 7. *If the rate of success on the reform λ_r increases, then*

1. *the effort on reform x_t increases;*
2. *the endogenous evaluation period Δ_t increases;*
3. *the voter's belief q_t decreases.*

An increase in the voter's payoff if the politician is fired λ_p is illustrated in Figure 7. As discussed in Section 2, an increase in λ_p is equivalently thought of as a decrease in the cost C of firing the incumbent and replacing him with a new politician believed to be competent with probability q_0 . This can be thought of as an institution that makes it easier to fire the incumbent. Panel (a) shows that the politician experiments more with the reform, Panel (b) shows that the voter gives the politician less time to implement the reform and Panel

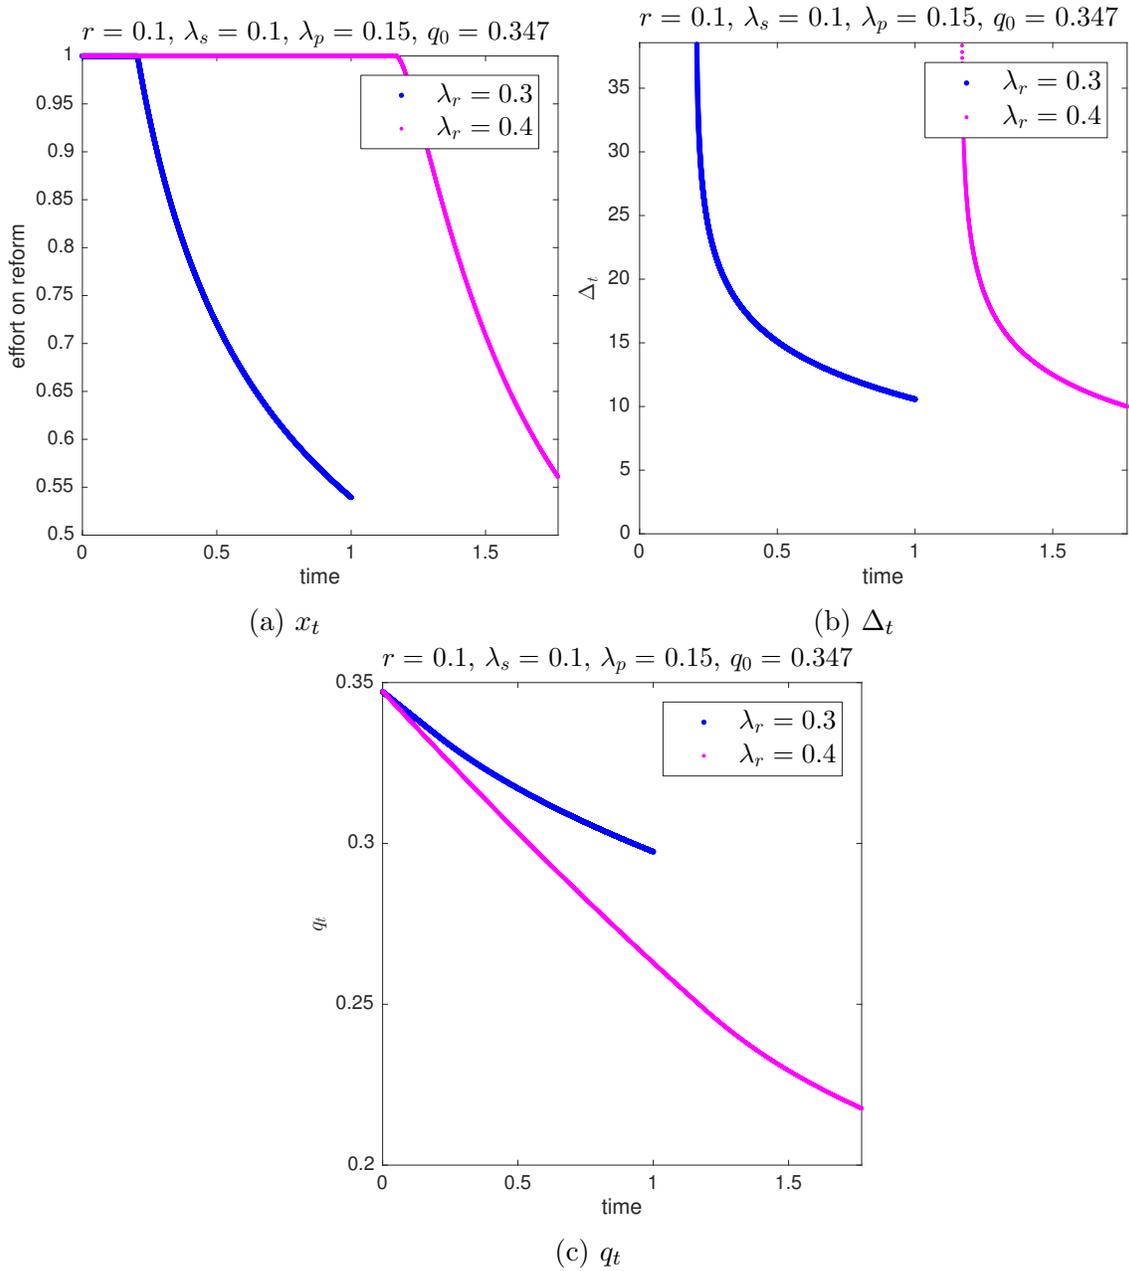


Figure 6: Increasing λ_r

(c) shows that the belief that the politician is good decreases at a faster rate with a higher value of λ_p . These changes are consistent with the intuition that the politician will have a greater incentive to exert effort on reform if he believes that he is more likely to be fired. Correspondingly, the voter gives him less time to achieve successes, because her incentive to fire is higher. With a lower cost of firing, there is more experimentation and hence more

efficient experimentation. In this example, total experimentation before T_1 increases from 0.76 to 0.87, and the efficient level of experimentation is 2.52. We summarize these results in Lemma 8.

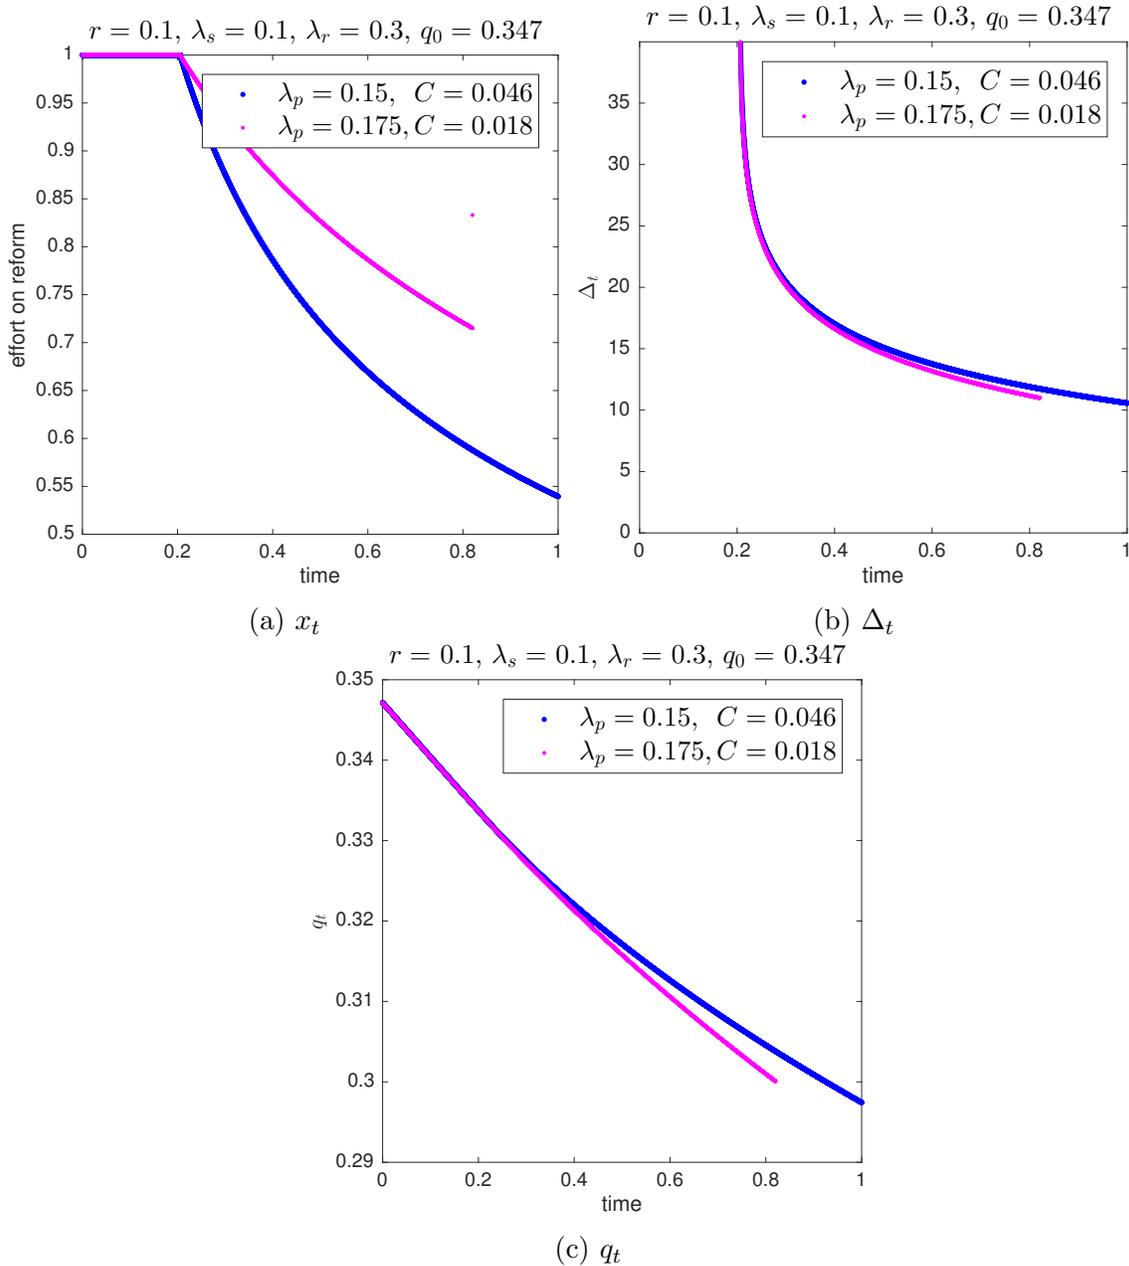


Figure 7: Increasing λ_p

Lemma 8. *If the return from firing the politician λ_p increases (equivalently, if the cost of firing the politician C decreases), then*

1. *the effort on reform x_t increases;*
2. *the endogenous evaluation period Δ_t decreases;*
3. *the voter's belief q_t decreases.*

5 Commitment

For comparison, we discuss the case in which the voter can commit to the criteria for firing. One can think of this case as one in which the voter can design an optimal firing policy for the politician, assuming that the politician will best respond to it at every instant given his information. The voter chooses the firing policy so as to maximize her payoff.

As before in the no-commitment case, the politician decides at every instant how much to work on the reform (devoting the remaining work resource to the status quo) as a function of his information. The voter decides at the outset the number of successes $N \geq 0$ needed for the politician to be kept, and then subsequently at every instant decides to fire the politician or not as a function of her information. If $N = \infty$, the voter chooses to never commit to retaining the politician.

5.1 One success with commitment

In the case of $N = 1$, the voter can commit to a time T to fire the unsuccessful politician. Given that the voter can commit to T , the politician will put full effort on reform at all times when $q_t > \lambda_s/\lambda_r$ and will put full effort on the status quo when $q_t < \lambda_s/\lambda_r$. The voter will anticipate this, and will want to hire the politician only up until the point where the unsuccessful politician reaches $q_t = \lambda_s/\lambda_r$, then fire the unsuccessful politician. Thus T is the time at which the politician (and voter's) belief reaches the myopic threshold T_m . From before, this is the solution to

$$\frac{q_0 e^{-\lambda_r t}}{q_0 e^{-\lambda_r t} + 1 - q_0} = \frac{\lambda_s}{\lambda_r}.$$

The voter hires the politician until T_m . If the politician gets a success before that, then he is given tenure. Otherwise, he is fired at time T_m . Note that in the case with only one success, the politician is fired sooner than is optimal, i.e. $T_m < T^*$ and there is under-experimentation.

5.2 Two successes with commitment

We consider the case when two successes are required for the politician to be retained, i.e. $N = 2$. We have the following result.

Proposition 3. *Consider the case with commitment. For $N = 2$, there exists a unique Markov equilibrium in which the politician is fired on the equilibrium path with some probability, and this increases experimentation relative to the case with no commitment.*

The proof of the first part again follows from the construction of the equilibrium. In this equilibrium, the politician puts full effort on reform all the time he is hired.²¹ The politician who has not gotten a success by a time \hat{T} (decided optimally by the voter) is fired. If the politician gets a success at time $t \leq \hat{T}$, the politician is kept for an additional duration Δ_t , also optimally decided by the voter. If the politician obtains no second success during that period, the politician is fired. Otherwise, the politician is hired forever.

Before \hat{T} . The politician plays action 1, hence his belief at $t \leq \hat{T}$ is

$$q_t = \frac{q_0 e^{-\lambda_r t}}{q_0 e^{-\lambda_r t} + 1 - q_0}.$$

Once the politician has obtained a success on the reform, he will put full effort on reform all the time, no matter the Δ_t decided by the voter. The voter prefers to keep a politician who is competent, and hence, Δ_t is the maximum possible duration that induces the politician to put full effort on reform before the first success. If $q_t \geq \lambda_s / \lambda_r$ (and so $t \leq T_m$), then there is no conflict of interest and the voter can set $\Delta_t = \infty$. If $t > T_m$, then the value Δ_t is chosen to make the politician indifferent between the reform and the status quo, so it solves

$$\lambda_r q_t \left[\frac{1}{r} - \frac{1}{r} e^{-(\lambda_r + r)\Delta_t} \right] - \lambda_s \left[\frac{1}{r} - \frac{1}{r} e^{-(\lambda_s + r)\Delta_t} \right] = (\lambda_r q_t - \lambda_s) W_t$$

where W_t is the politician's continuation value at t if he has not obtained any success by that time.

As in the case of no commitment, we calculate the equilibrium belief and endogenous term at the time \hat{T} . We have the following result.

Lemma 9. *The values $\Delta_{\hat{T}}$ and $q_{\hat{T}}$ are given implicitly by*

$$(\lambda_r q_{\hat{T}} + r) \frac{\lambda_p}{r} = \lambda_r q_{\hat{T}} [1 + V^R(\Delta_{\hat{T}})],$$

²¹The politician is, in fact, indifferent along the equilibrium no-success path, so he could put any amount of effort into reform. For this reason, he has no incentive to deviate.

$$V^R(\Delta_{\hat{T}}) = \frac{\lambda_r}{r} - \left[\frac{\lambda_r}{r} - \frac{\lambda_p}{r} \right] e^{-(r+\lambda_r)\Delta_{\hat{T}}}$$

and

$$\lambda_r q_{\hat{T}} \left[\frac{1}{r} - \frac{1}{r} e^{-(\lambda_r+r)\Delta_{\hat{T}}} \right] - \lambda_s \left[\frac{1}{r} - \frac{1}{r} e^{-(\lambda_s+r)\Delta_{\hat{T}}} \right] = 0.$$

Lemma 9 gives the equations that determine the values of $\Delta_{\hat{T}}$, $q_{\hat{T}}$ and \hat{T} . We use these as the initial conditions to determine the equilibrium values at $t < \hat{T}$.

As before we solve this system of equations numerically and obtain results as illustrated below in Figure 8. Figure 8 panel (a) shows that beliefs fall faster in the case of commitment, because more effort is being exerted on the reform. Experimentation is thus higher with commitment and closer to the first best. Figure 8 panel (b) illustrates that the voter gives the politician the same amount of time to get the first success with commitment, but gives less time for the second success, once a first success is achieved. With commitment, the voter can enforce this shorter term.

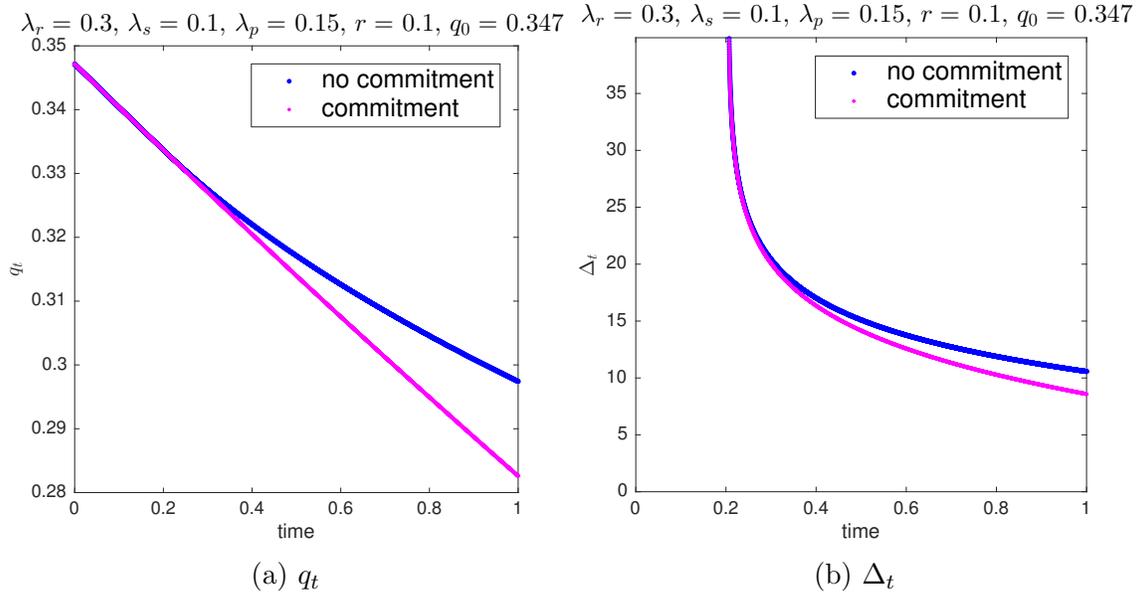


Figure 8: Equilibrium values assuming no success. No commitment versus commitment

As we do not observe sharp decreases in experimentation in practice, we believe the no commitment case is the empirically relevant one, and we take this to the data in the next section.

6 Empirical evidence

In this section, use data on financial reforms to see if we can observe evidence of reform fatigue in countries with presidential political systems. While we do not view this simple empirical exercise to be a direct test of the theory, we believe it supports the model’s empirical validity. From the stylized model, we interpret T_1 as the first election after the politician has held office for some time, when he could be potentially fired. The time T_1 is endogenous in the model. However elections in presidential systems typically occur after a fixed number of years as determined by the country’s constitution, and thus the time of the election is exogenous in our sample.²² We show that, despite this disjuncture, the empirical patterns for these systems are consistent with the predictions of the model.

We focus on presidential countries since the theory is most applicable to these regimes. Under presidential systems, heads of government are directly elected by voters. In contrast, under parliamentary systems, heads of government are elected by legislators, who tend to have more information than voters about the activities of the leader, and, in particular, about reform efforts exerted by leaders. Since a key ingredient of the model is that voters cannot directly observe the leader’s effort on reform, we conclude that reform fatigue as described in the theory should not manifest in parliamentary systems. We empirically verify this claim below.

At first glance, varying term lengths under parliamentary systems may seem to bear closer resemblance to the endogenous term length in the model. However, it is worth noting that the way in which term lengths are endogenous under a parliamentary system differs from the way in which term lengths are endogenous under the model. Under parliamentary systems, the decision to call an election at a particular time typically occurs ex-post, after the start of the term, and in response to the unfolding political and economic conditions in the country. In contrast, “term lengths” in the model are decided ex-ante, and fixed just prior to the start of that term. For example, the length of the first term, denoted T_1 on Figure 2, is a function of factors such as voters’ prior beliefs regarding the politicians’ ability to pursue reforms and the status-quo. The length of the second term, $T_2 - T_1$ in Figure 2, is also decided before the end of the first term, and adjusted as a function of the politician’s performance during the first term (hence prior to the second term). As such, performance in the second term impacts voters’ beliefs about the politician, but not the timing of firing.

The model generates two key predictions with regard to the pattern of reform effort in

²²Since some elections do deviate from their pre-determined timing, for example as a result of a coup, we also examine whether the timing of elections influences our results when presenting the empirical results below.

the period up to T_1 . First, as demonstrated in Figure 3, in phase I (before the first success), there is a general pattern of reform fatigue. From the time the politician takes office to the end of his term at T_1 , reform effort is non-increasing, even if the pace of decline may vary. Second, as illustrated in Figure 4, in phase II (after first success, before second success), we see that there are two cases, which hold different implications for how output affects the reform cycle. In the first case, if a success (high output) is observed before T_1 , and this success was achieved on the reform, the politician will exert full effort on the reform thereafter. This case suggests high output will mitigate reform fatigue, or, increase subsequent effort on reforms. In contrast, if a success is observed before T_1 , but was achieved on the status quo, then the politician will exert full effort on the status quo. This second case implies that high output will reinforce reform fatigue, or, decrease subsequent effort on reform. In short, under the theory, high output from reform will mitigate reform fatigue while high output from the status quo will exacerbate it. Although we cannot directly observe if output stems from reform, we can examine whether the empirical reform patterns differ based on output changes and the co-occurrence of output changes and reforms during the previous period.

We focus on financial reforms for three reasons. First, much of the qualitative debate and anecdotal accounts of reform fatigue have focused on reforms within the financial sector. Second, financial reforms are typically implemented by an executive and implementation requires effort to design the policy, and to build coalitions that will support the relevant legislation and write the text of this legislation. This effort of the politician is arguably unobserved to the voter, as it is complicated to implement, requiring reasonably sophisticated legislation and execution. Third, our theory assumes that effort on the reform arm has a larger probability of generating output than the status quo arm. This assumption is supported by past work which has provided empirical evidence around the positive economic impact of financial reforms (e.g., Prati et al., 2013; Christiansen et al., 2013).²³

6.1 Data

We examine political cycles in financial reforms using cross-country panel data. Our reform measure is a market liberalization index for the financial sector from Abiad et al. (2008), which covers the 1973-2005 period. The index aggregates seven different aspects of financial reforms, and we rescale it to lie between 0 and 100. (Appendix B provides further details on this variable and other variables used in the analysis). The change in the financial liberalization

²³One might consider reforms of other economic policies, such as agriculture or trade. We do not focus on these either because there is insufficient variation in reform episodes (as in the case of agriculture), or because they are arguably observable to the voter (as with trade policy).

index, as measured by the annual first difference, is our main outcome variable – as this captures the extent of financial reform.

We use data on national executive and legislative elections from the World Bank Database of Political Institutions (DPI), which covers the 1975-2012 period. Our sample is comprised of countries with presidential systems, in which the head of the executive branch is elected either directly or by an electoral college.

Finally, to capture the binary nature of output generated by reforms under the model, we construct an indicator variable for whether the growth rate of GDP per capita is above trend or not. Specifically, for each country, we use the Hodrick and Prescott (1981, 1997) (HP) filter to extract the trend component of GDP per capita, with a smoothing parameter of 400.²⁴ The variable *Output indicator* is defined to be one if lagged GDP per capita is above trend, i.e., if the output gap is positive. We also use a measure of whether the country participates in the IMF program, which is used for additional analysis in the appendix.

Our regression sample covers 56 countries between years 1976 and 2004. Appendix Table B.1 provides summary statistics for key variables.

6.2 Empirical results

To examine whether annual changes in financial reforms vary over the course of the electoral cycle, we estimate:

$$\Delta Ref_{ct} = \beta_1(Lag)_{ct} + \beta_2(Year\ of)_{ct} + \beta_3(Lead)_{ct} + \varphi_c + \delta_t + \varepsilon_{ct}, \quad (14)$$

where ΔRef_{ct} is the first difference in the financial liberalization index, Ref_{ct} , for a given country c and for a year t . φ_c are country fixed effects which sweep out time invariant country characteristics and δ_t are year fixed effects which control for financial reform changes common to the global sample. The variable $(Year\ of)_{ct}$ is equal to one if the country has an election in that year; Lag_{ct} denotes the year after the election; and $Lead_{ct}$ denotes the year before an election.²⁵ Based on the model’s prediction, an increase in the pace of reforms after an election corresponds to a positive sign on β_1 , while a slowdown in the pace of reforms prior

²⁴The use of 400 as a smoothing parameter follows, for instance, Cooley and Ohanian (1991) and Correia et al. (1992). However, in the empirical appendix (Appendix Table B.2), we also verify that our results are robust to using a smoothing parameter of 100, which has been used, for example, by Backus and Kehoe (1992) and Barro and Ursúa (2008). We focus on these two parameters since both have been commonly used (Ravn and Uhlig, 2002).

²⁵Abiad and Mody (2005) also estimate an effect of the *Lag* variable on the change in reforms, which they reference as the first year or “honeymoon” effect. In contrast, our specification traces effects over the entire election cycle by separately estimating the impact of the *Lag*, *Lead* and *(Year of)* election variables.

Table 1: Financial reforms in presidential regimes

	Executive or Legislative			Executive			Legislative		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
<i>Lag</i>	0.748** (0.374)	0.915** (0.383)		0.984** (0.459)	1.157*** (0.441)		0.709* (0.429)	0.952** (0.438)	
<i>Year of</i>	0.051 (0.358)	0.291 (0.328)	-0.199 (0.308)	0.005 (0.409)	0.185 (0.381)	-0.299 (0.370)	0.059 (0.417)	0.386 (0.368)	-0.195 (0.358)
<i>Lead</i>	-0.670* (0.378)		-0.849** (0.376)	-0.568 (0.459)		-0.862** (0.439)	-0.875** (0.400)		-1.060*** (0.399)
R-squared	0.154	0.151	0.151	0.153	0.152	0.150	0.155	0.151	0.152
Observations	1195	1195	1195	1195	1195	1195	1195	1195	1195
No. countries	56	56	56	56	56	56	56	56	56

Notes: The dependent variable is the change in the financial reform index. All regressions include country and year fixed effects. Standard errors are in parentheses, clustered at the country level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level.

to an election corresponds to a negative sign on β_3 . We estimate equation (14) using OLS, and cluster the standard errors at the country level to account for potential serial correlation over time.

Table 1 presents estimates of the main specification. The first three columns consider the impact of executive and legislative elections together; the second three columns consider just executive elections; and the last three columns consider just legislative elections. Within each election type, we first present estimates of equation (14), which simultaneously includes indicators for the lead, the lag, and the year of election variables. In these specifications, the omitted category is the year(s) in between the lag and lead years. Since average term lengths are three years long, this suggests we may have limited power to identify all three effects. Thus we also present additional estimates looking separately at just the election lead or election lag, alongside the year of election variable. The omitted category in these specifications changes to all years in the term not included as regressors.

The results in Table 1 present evidence consistent with a political cycle in financial reforms. The positive coefficient on the *Lag* variable indicates that reforms are implemented to a greater degree right after elections, while the negative coefficient on the *Lead* variable suggests that the implementation of reforms diminish in the run up to the next election. The implied effects are substantial. Consider column (1). The coefficient of 0.748 on *Lag* implies that after an election, reforms increased by 36 percent relative to the mean change of 2.099. The coefficient of -0.670 on *Lead* implies that reform implementation slowed by 32 percent in the year before an election. These estimates document substantial differences in the extent of financial market liberalization at the beginning and end of a politician's term. This pattern is even stronger in the other columns when either the *Lag* or *Lead* term is omitted. The

disaggregation of the legislative and executive elections also indicate that the effects are not driven by either type of election, as the coefficients on *Lead* and *Lag* across specifications are not statistically distinguishable from one another at the 5 percent level.

In Appendix B, we present several additional results. First, although most presidential elections are pre-determined by constitutionally mandated term lengths, some term lengths deviate from this mandate, which raises concerns that the late exit of a current executive or early entry of a new executive may reflect economic conditions, including financial reforms. We rule out that this type of reverse causality drives our results by drawing on an approach used by Brender and Drazen (2005). We show that the results for executive elections remain in place even when we restrict the sample to terms that did not deviate from their constitutionally pre-determined term length (Appendix Table B.3, columns 1-3).

We also show in Appendix B that the patterns in the data are consistent with other results from the theory. We find stronger evidence of reform fatigue during an executive's first term in office (Appendix Table B.3, columns 4-9), which is when the model predicts learning effects are likely to be strongest. Looking across election types, we observe no evidence that reform cycles are significantly larger (or smaller) in countries that participate in IMF programs (Appendix Table B.4). This suggests that patterns of fatigue are not driven by external pressure from this international organization. Finally, we observe no evidence of reform cycles within parliamentary regimes (Appendix Table B.5). The differing patterns across regime types is consistent with the idea that executives in parliamentary systems are not directly elected by voters.

Finally, we examine if these patterns of reform fatigue vary based on the occurrence of output changes, as predicted by the model and illustrated in Figure 4. As discussed above, under the theory, high output can exacerbate the slow-down in reforms (if the rise in output stems from the status quo); or it can mitigate this slow-down and facilitate continued reform (if the output stems from reforms). First, in Panel A of Table 2, we examine which effect tends to hold by interacting the election lead with *Output indicator*, which equals one if (lagged) output was above trend. The significant positive coefficients on this interaction term indicate that relatively high output tends to mitigate reform fatigue. Under the theory, this dampening in the reform cycle is consistent with the attribution of initial success and associated output to reforms.

Second, although positive output may also be generated from the status quo, we would expect this to be less probable if reform effort is high close to the end of the term. To investigate if the co-occurrence of high output and reforms mitigate reform fatigue, in Panel B of Table 2, we utilize a more refined specification that interacts the election lead with an indicator which equals one if output was above trend *and* non-negative reforms occurred

Table 2: Financial reforms and output changes

	Panel A: Effects by lag output changes					
	Executive or Legislative		Executive		Legislative	
	(1)	(2)	(3)	(4)	(5)	(6)
Lag	0.776**		1.154**		0.672	
	(0.359)		(0.455)		(0.423)	
Year of	0.099	-0.165	0.059	-0.305	0.145	-0.095
	(0.368)	(0.321)	(0.388)	(0.355)	(0.420)	(0.353)
Lead	-1.252***	-1.423***	-1.181**	-1.513***	-1.145***	-1.319***
	(0.406)	(0.404)	(0.507)	(0.503)	(0.438)	(0.442)
Output indicator	-1.371***	-1.346***	-1.282***	-1.234***	-1.187***	-1.184***
	(0.393)	(0.393)	(0.371)	(0.375)	(0.391)	(0.392)
Lead×Output indicator	1.284**	1.259**	1.361**	1.332*	0.812	0.814
	(0.518)	(0.516)	(0.688)	(0.680)	(0.583)	(0.580)
R-squared	0.163	0.160	0.164	0.159	0.161	0.159
Observations	1189	1189	1189	1189	1189	1189
No. countries	56	56	56	56	56	56
	Panel B: Effects by lag output and reform changes					
	Executive or Legislative		Executive		Legislative	
	(1)	(2)	(3)	(4)	(5)	(6)
Lag	0.742**		1.110**		0.659	
	(0.355)		(0.453)		(0.418)	
Year of	0.078	-0.174	0.051	-0.299	0.119	-0.116
	(0.366)	(0.321)	(0.387)	(0.354)	(0.418)	(0.355)
Lead	-1.330***	-1.499***	-1.320**	-1.648***	-1.366***	-1.535***
	(0.473)	(0.460)	(0.638)	(0.613)	(0.488)	(0.475)
Output and reforms indicator	-1.297***	-1.279***	-1.199***	-1.170***	-1.163***	-1.160***
	(0.377)	(0.379)	(0.351)	(0.359)	(0.364)	(0.365)
Lead×Output and reforms indicator	1.566***	1.554***	1.816**	1.806**	1.363**	1.361**
	(0.593)	(0.589)	(0.849)	(0.842)	(0.577)	(0.571)
R-squared	0.160	0.157	0.161	0.156	0.158	0.156
Observations	1187	1187	1187	1187	1187	1187
No. countries	56	56	56	56	56	56

Notes: The dependent variable is the change in the financial reform index. *Output indicator* is an indicator variable equal to 1 when lagged output per capita is greater than its lagged trend value. *Output and reforms indicator* is an indicator variable equal to 1 when *Output indicator* is 1 and the lagged change in financial reform index is greater than or equal to zero. All regressions include country and year fixed effects. Standard errors are in parentheses, clustered at the country level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level.

during the last period.²⁶ Here again, we observe significant, positive coefficients on the interaction terms, suggesting that high output tends to counter the slow-down in reforms observed at the end of the electoral term.²⁷ These results are consistent with the idea that

²⁶We focus on non-negative reforms, rather than strictly positive reforms, because it may take effort to simply maintain financial liberalization. This is underscored by the observation that a substantial portion of the cases of zero changes in financial reforms occurs in countries that already have a relatively high level of financial liberalization. Specifically, one-third of the zero changes in our sample are from countries that have a financial reform index of at least 61.9, which lies at the 67th percentile of the index of financial reforms.

²⁷Appendix Table B.2 also shows that these results look similar if we define above-trend output using an

politicians will continue pursuing reforms, in the presence of initial success on reforms.

7 Conclusion

This paper presents a rational theory of reform fatigue. The theory is based on the voter's uncertainty about the competence of the politician. Success on reforms reflect a competent politician, whereas failure may result in the politician losing office. In an equilibrium with reform fatigue, the voter is unable to know the politician's type with certainty when there is a conflict of interest. For this to happen, the politician must exert an intermediate amount of effort on the reform, so that if a success occurs, it cannot be attributed with certainty to the reform or the status quo. As a result, the voter's pace of learning about the politician's type may vary from the politician's.

For an intermediate amount of effort to be allocated to the reform in equilibrium, the politician must be indifferent between the reform and the status quo as long as there is a conflict of interest. The voter gives the politician an evaluation period that depends on the timing of the first success to create this indifference. The politician's belief that he is competent decreases over time with no success, while his time horizon keeps shrinking. This increases the politician's incentive to switch to the status quo. In order to balance the politician's changes in beliefs and payoffs, the voter decreases the endogenous term length as long as there is no success. The decrease in term length decreases the benefit of a success, and hence decreases the incentive to switch to the status quo. The model predicts that if the reform generates a success, then we should observed an increase in reforms. Otherwise, we should see a decrease in reforms.

We present empirical evidence consistent with the predictions of the model. A reform fatigue cycle is identified in financial reform data across countries with presidential systems. Our theory suggests that these reform cycles are due to politicians optimally choosing to experiment with reforms under the shadow of electoral incentives. Our empirical results corroborate this prediction by showing a positive correlation between output and reforms as predicted by the model when the politician is competent.

The evidence we provide suggests that the model is empirically valid, but we leave a more direct test for future work. Such a test might examine other types of reforms in which effort is unobservable but the evaluation term is endogenous, and focus on a context where it is possible to directly discern if success stems from pulling one policy lever over another.

alternate smoothing parameter of 100.

A Theoretical appendix

A.1 Proof of Lemma 1

Substituting p_0 from equation (2) into the expression for p_t in equation (3) gives

$$p_t = \left[1 + e^{(\lambda_r - \lambda_s)t} \frac{\lambda_s(1 - x_\tau)}{\lambda_r x_\tau q_\tau} \right]^{-1}. \quad (15)$$

The voter will fire the politician only at time $\tau + \Delta_\tau$ when he is indifferent between firing the politician and keeping him one more instant. If the voter keeps the politician one more instant, with probability $p_{\Delta_\tau} \lambda_r$ there is a success on the reform and the politician puts full effort on reform thereafter, giving the voter a payoff $[1 + \frac{\lambda_r}{r}]$. With probability $(1 - p_{\Delta_\tau}) \lambda_s$ there is a success on the status quo and the politician, who is still uncertain of his competence, pursues the optimal strategy given his belief q_τ and delivers the payoff $[1 + H(q_\tau)]$ to the voter. If there is no success, the voter strictly prefers to fire the politician and obtains the payoff $\frac{\lambda_p}{r}$. Thus the voter's indifference condition at time $\tau + \Delta_\tau$ is

$$\begin{aligned} p_{\Delta_\tau} \lambda_r \left[1 + \frac{\lambda_r}{r} \right] dt + (1 - p_{\Delta_\tau}) \lambda_s [1 + H(q_\tau)] dt \\ + [1 - p_{\Delta_\tau} \lambda_r dt - (1 - p_{\Delta_\tau}) \lambda_s dt - r dt] \frac{\lambda_p}{r} = \frac{\lambda_p}{r}. \end{aligned}$$

Rearranging the voter's indifference condition gives

$$p_{\Delta_\tau} = \left[1 - \frac{(\lambda_r - \lambda_p)(1 + \lambda_r/r)}{\lambda_s(1 + H(q_\tau)) - \lambda_p(1 + \lambda_s/r)} \right]^{-1} = \left[1 + e^{(\lambda_r - \lambda_s)\Delta_\tau} \frac{\lambda_s(1 - x_\tau)}{\lambda_r x_\tau q_\tau} \right]^{-1}.$$

The last equality follows from equation (15), which gives the voter's belief p_t at any time t after τ . This is evaluated at $t = \Delta_\tau$. Rearranging the last equality gives

$$x_\tau = \left[1 - e^{-(\lambda_r - \lambda_s)\Delta_\tau} \frac{\lambda_r q_\tau (\lambda_r - \lambda_p) (1 + \lambda_r/r)}{\lambda_s [\lambda_s(1 + H(q_\tau)) - \lambda_p(1 + \lambda_s/r)]} \right]^{-1}. \quad (16)$$

■

A.2 Proof of Lemma 2

After a first success on the reform at time τ , the politician exerts full effort on the reform in phase II of the equilibrium. If another success is obtained before time $\tau + \Delta_\tau$, then the

politician moves to phase III of the equilibrium in which he is kept forever and maintains full effort on the reform because he knows that he is good at that time. The voter's payoff in phase III is thus $\frac{\lambda_r}{r}$. The voter's payoff after a first success on the reform is thus

$$V_t^R = \lambda_r dt \left[1 + \frac{\lambda_r}{r} \right] + (1 - \lambda_r dt - r dt) V_{t+dt}^R.$$

Simplifying, gives the ODE for V_t^R , which is $-\frac{dV_t^R}{dt} = \lambda_r \left[1 + \frac{\lambda_r}{r} \right] - (\lambda_r + r) V_t^R$. The voter fires the politician and time Δ_τ , and thus the boundary condition is $V_{\Delta_\tau}^R = \lambda_p/r$. Solving this ODE gives the continuation payoff for the voter if the politician got a first success on the reform and time τ . This is

$$V_\tau^R = \frac{\lambda_r}{r} - \left[\frac{\lambda_r}{r} - \frac{\lambda_p}{r} \right] e^{-(r+\lambda_r)\Delta_\tau}. \quad (17)$$

After a first success on the status quo at time τ , the politician exerts full effort on the status quo in phase II of the equilibrium. As before, if another success is obtained before time Δ_τ , then the politician moves to phase III of the equilibrium in which he is kept forever and does the optimal experimentation for the voter, given that his belief that he is competent is q_τ . The voter's payoff in phase III is thus $H(q_\tau)$. The voter's payoff after a first success on the status quo is thus

$$V_t^S = \lambda_s dt [1 + H(q_\tau)] + (1 - \lambda_s dt - r dt) V_{t+dt}^S$$

Simplifying gives $-\frac{dV_t^S}{dt} = \lambda_s [1 + H(q_\tau)] - (\lambda_s + r) V_t^S$ with boundary condition $V_{\Delta_\tau} = \lambda_p/r$. The continuation payoff for the voter at τ if the politician got a first success on the status quo at τ is thus

$$V_\tau^S = \frac{1 + H(q_\tau)}{1 + r/\lambda_s} - \left[\frac{1 + H(q_\tau)}{1 + r/\lambda_s} - \frac{\lambda_p}{r} \right] e^{-(r+\lambda_s)\Delta_\tau}. \quad (18)$$

■

A.3 Proof of Lemma 3

Recall W_t is the continuation payoff of the politician, at time t , who has not obtained any success by t . We get that W_t is given by

$$W_t = x_t q_t \lambda_r dt W_t^R + (1 - x_t) \lambda_s dt W_t^S + [1 - x_t q_t \lambda_r dt - (1 - x_t) \lambda_s dt - r dt] W_{t+dt},$$

where W_t^R and W_t^S are the politician's continuation payoffs after a success on the reform and status quo respectively. Since the politician is indifferent between the reform and the status quo, we can set $x_t = 0$ in the above expression. Using the approximation that $W_{t+dt} = W_t + dW_t$ we have that W_t evolves according to the ODE

$$\frac{dW_t}{dt} = W_t(\lambda_s + r) - \lambda_s W_t^S. \quad (19)$$

We calculate W_t^S . If the politician achieves a success on the status quo at time t , then he has Δ_t units of time to obtain the second success. The politician will put full effort on the status quo during this time. If the second success is obtained before Δ_t , then the politician is retained permanently and receives discounted payoff $\frac{e^{-r\Delta_t}}{r}$. The probability that at least one success is obtained in the interval $(t, t + \Delta_t]$ is $1 - e^{-\lambda_s \Delta_t}$. The politician receives the payoff $\frac{1}{r} [1 - e^{-r\Delta_t}]$ between t and Δ_t , and thus the politician's payoff after a first success on the status quo is

$$W_t^S = \frac{e^{-r\Delta_t}}{r} [1 - e^{-\lambda_s \Delta_t}] + \frac{1}{r} [1 - e^{-r\Delta_t}] = \frac{1}{r} [1 - e^{-(\lambda_s + r)\Delta_t}].$$

Substituting into equation (19) gives

$$\frac{dW_t}{dt} = W_t(\lambda_s + r) - \frac{\lambda_s}{r} [1 - e^{-(\lambda_s + r)\Delta_t}].$$

We thus have a differential equation for W_t with boundary condition $W_{T_1} = 0$ since the politician is fired at time T_1 if there is no success. We obtain the closed form solution of W_t

$$W_t = e^{(\lambda_s + r)t} \int_{T_1}^t \frac{\lambda_s [e^{-(\lambda_s + r)(z + \Delta_z)} - e^{-(\lambda_s + r)z}]}{r} dz. \quad (20)$$

■

A.4 Proof of Lemma 4

We get q_{T_1} as a function of Δ_{T_1} by solving the indifference condition for the politician. Noting that $W_{T_1} = 0$, we have

$$\lambda_r q_{T_1} [1 - e^{-(r+\lambda_r)\Delta_{T_1}}] = \lambda_s [1 - e^{-(r+\lambda_s)\Delta_{T_1}}],$$

which gives

$$q_{T_1} = \frac{\lambda_s [1 - e^{-(r+\lambda_s)\Delta_{T_1}}]}{\lambda_r [1 - e^{-(r+\lambda_r)\Delta_{T_1}}]}.$$

We get x_{T_1} as a function of Δ_{T_1} by solving the indifference condition for the voter, to fire the politician at T_1 or to wait one instant later. This indifference condition is

$$\lambda_r x_{T_1} q_{T_1} (1 + V_{T_1}^R) + \lambda_s (1 - x_{T_1}) (1 + V_{T_1}^S) - (r + \lambda_s (1 - x_{T_1}) + \lambda_r x_{T_1} q_{T_1}) \frac{\lambda_p}{r} = 0,$$

where $V_{T_1}^R$ and $V_{T_1}^S$ are the voter's continuation payoff after a success on the reform and status quo respectively given by equations (17) and (18) respectively. We get

$$x_{T_1} = \frac{\lambda_p - \lambda_s (1 + V_{T_1}^S - \lambda_p/r)}{\lambda_r q_{T_1} (1 + V_{T_1}^R - \lambda_p/r) - \lambda_s (1 + V_{T_1}^S - \lambda_p/r)}.$$

We get Δ_{T_1} by solving the indifference condition for the voter, to fire the politician at $T_1 + \Delta_{T_1}$ or to wait one instant later:

$$\left[1 - \frac{\lambda_r (1 + \lambda_r/r) - \lambda_p (1 + \lambda_r/r)}{\lambda_s (1 + H(q_{T_1})) - \lambda_p (1 + \lambda_s/r)} \right]^{-1} = \left[1 + e^{(\lambda_r - \lambda_s)\Delta_{T_1}} \frac{\lambda_s (1 - x_{T_1})}{\lambda_r x_{T_1} q_{T_1}} \right]^{-1}.$$

■

A.5 Proof of Lemma 6

(To be included)

A.6 Proof of Lemma 7

(To be included)

A.7 Proof of Lemma 8

(To be included)

A.8 Proof of Lemma 9

The continuation value of the unsuccessful politician W_t evolves according to the ordinary differential equation (ODE)

$$\frac{dW_t}{dt} - (\lambda_s + r)W_t = -\frac{\lambda_s}{r} [1 - e^{-(r+\lambda_s)\Delta(q_t, W_t)}],$$

with boundary condition $W_{\hat{T}} = 0$.

The voter should be indifferent between firing the politician at \hat{T} , or waiting one instant later. This indifference condition is

$$(\lambda_r q_{\hat{T}} + r) \frac{\lambda_p}{r} = \lambda_r q_{\hat{T}_m} [1 + V^R(\Delta_{\hat{T}})]$$

where $V^R(\Delta_{\hat{T}})$ is the voter's continuation value at time \hat{T} right after the politician got a success from the reform, it can be expressed in closed form

$$V^R(\Delta_{\hat{T}}) = \frac{\lambda_r}{r} - \left[\frac{\lambda_r}{r} - \frac{\lambda_p}{r} \right] e^{-(r+\lambda_r)\Delta_{\hat{T}}}$$

and we also have $q_{\hat{T}}$ as a function of $\Delta_{\hat{T}}$ from the indifference condition of the politician:

$$\lambda_r q_{\hat{T}} \left[\frac{1}{r} - \frac{1}{r} e^{-(\lambda_r+r)\Delta_{\hat{T}}} \right] - \lambda_s \left[\frac{1}{r} - \frac{1}{r} e^{-(\lambda_s+r)\Delta_{\hat{T}}} \right] = 0.$$

■

B Empirical appendix

B.1 Data

The financial market reform index from Abiad et al. (2008) is comprised of seven different aspects: (i) credit controls and excessively high reserve requirements (including directed credit and credit ceilings), (ii) interest rate controls, (iii) entry barriers, (iv) state ownership in the banking sector, (v) capital account restrictions, (vi) prudential regulations and supervision of the banking sector, and (vii) securities market policy. For each aspect of liberalization, the country is assigned a score ranging from 0 to 3, with zero being fully repressed and three being fully liberalized. An aggregate index is constructed by summing up all the categories. We rescale this index to be between 0 and 100.

Our primary source for national elections data is the World Bank Database of Political Institutions (DPI), which covers 178 countries over the 1975-2012 period. Other sources for elections data with large coverage include the Institutions and Elections Project (IAEP) and Golder (2005). We detected some inconsistencies between DPI and these two databases; hence, corrections were made to address them. We make the following changes: no executive elections in Madagascar, 1977, and Mexico, 1997; executive and legislative elections in Colombia, 1998 instead of 1999, and in Kenya, 1988 instead of 1987; executive election in Zimbabwe in 1990.

For data on GDP per capita, we utilize the variable “rgdpch” from the Penn World Tables 7.0, defined as PPP converted GDP per capita (chain series), at 2005 constant prices. Appendix Table B.1 provides summary statistics for our key variables.

Country sample

Our sample consists of 56 countries with presidential election systems:

Africa: Algeria, Burkina-Faso, Cameroon, Cote d’Ivoire, Ethiopia, Ghana, Israel, Jordan, Kenya, Madagascar, Morocco, Mozambique, Nigeria, Senegal, Tanzania, Tunisia, Uganda, Zimbabwe.

Americas: Argentina, Bolivia, Brazil, Chile, Colombia, Costa Rica, Dominican Republic, Ecuador, El Salvador, Guatemala, Mexico, Nicaragua, Paraguay, Peru, United States, Uruguay, Venezuela.

Asia: Bangladesh, South Korea, Nepal, Pakistan, Philippines, Sri Lanka, Taiwan, Thailand.

Europe & Central Asia: Azerbaijan, Belarus, Georgia, Kazakhstan, Kyrgyz Republic, Lithuania, Poland, Portugal, Russia, Spain, Turkey, Ukraine, Uzbekistan.

Table B.1: Summary statistics (1976-2004)

	Obs.	Mean	Std. dev.	Min.	Max.
Financial reform index	1195	43.42	27.36	0	100
Change in financial reform index	1195	2.099	5.365	-28.57	38.10
Executive or legislative election	1195	0.268	0.443	0	1
Executive election	1195	0.172	0.377	0	1
Legislative election	1195	0.223	0.416	0	1
Observed executive or legislative election term length	294	3.425	2.382	1	16
Observed executive election term length	175	4.823	1.718	1	16
Observed legislative election term length	234	4.141	2.017	1	16
Output indicator	1195	0.495	0.500	0	1
Output and reforms indicator	1195	0.465	0.499	0	1
IMF program participation	1195	0.549	0.498	0	1

Notes: This table presents summary statistics for our sample which comprises 56 countries with presidential regimes over 1976-2004. An observation for term length is one observed election cycle, i.e., election to election. Years up to first election observed in the sample are not counted as a term. *Output indicator* is an indicator variable equal to 1 when lagged output per capita is above or equal to its lagged trend value. *Output and reforms indicator* is an indicator variable equal to 1 when *Output indicator* is 1 and the lagged change in financial reform index is greater than or equal to zero. *IMF program participation* is equal to 1 if the country has an outstanding loan with the IMF.

B.2 Robustness checks

In Appendix Table B.2, we test whether the results in Table 2 are robust to an alternative construction of the trend for output per capita. Specifically, we change the smoothing parameter to the value of 100 for the HP filter. The results with this measure are also consistent with the idea that the pace of reforms slows down less when there is relatively high output growth, particularly when this high output growth occurs alongside non-negative reforms during the previous period.

In the Appendix Table B.3, we present and discuss additional results that are consistent with the implications of the model. Firstly, for the most part, presidential elections occur on the basis of constitutionally mandated term lengths. However, when elections deviate from the pre-determined schedule, it raises the potential concern that a change in government could have arisen, for example, in response to financial reforms or associated output changes. We rule out this type of reverse causality by restricting the sample to terms that were not deviants from the constitutionally mandated term length. For the 56 countries in our sample, we assemble data on constitutions over the sample period, and record the chief executive's term length as mandated by the constitution in effect for every year. We identify and remove cases in which the actual term length deviates from the pre-determined term length. This includes both cases when the constitutionally mandated term length does not change, but the actual term length is shorter or longer than the pre-determined interval (which would arise if an

Table B.2: Financial reforms and output changes (using HP filter with alternate parameter)

Panel A: Effects by lag output changes						
	Executive or Legislative		Executive		Legislative	
	(1)	(2)	(3)	(4)	(5)	(6)
Lag	0.742**		1.119**		0.648	
	(0.353)		(0.457)		(0.418)	
Year of	0.090	-0.161	0.051	-0.302	0.144	-0.087
	(0.364)	(0.317)	(0.386)	(0.352)	(0.417)	(0.352)
Lead	-1.147***	-1.317***	-1.240**	-1.556***	-1.046**	-1.223***
	(0.424)	(0.431)	(0.525)	(0.521)	(0.455)	(0.464)
Output indicator	-0.981**	-0.973**	-0.971***	-0.935**	-0.822**	-0.833**
	(0.382)	(0.382)	(0.369)	(0.371)	(0.380)	(0.379)
Lead×Output indicator	1.001*	0.992*	1.400*	1.361*	0.561	0.582
	(0.598)	(0.600)	(0.761)	(0.750)	(0.685)	(0.683)
R-squared	0.158	0.155	0.160	0.155	0.157	0.155
Observations	1189	1189	1189	1189	1189	1189
No. countries	56	56	56	56	56	56
Panel B: Effects by lag output and reform changes						
	Executive or Legislative		Executive		Legislative	
	(1)	(2)	(3)	(4)	(5)	(6)
Lag	0.712**		1.087**		0.634	
	(0.352)		(0.456)		(0.417)	
Year of	0.068	-0.174	0.042	-0.301	0.117	-0.108
	(0.363)	(0.317)	(0.386)	(0.352)	(0.418)	(0.354)
Lead	-1.217**	-1.385***	-1.354**	-1.670***	-1.260**	-1.432***
	(0.475)	(0.471)	(0.644)	(0.621)	(0.490)	(0.480)
Output and reforms indicator	-0.885**	-0.882**	-0.869**	-0.851**	-0.776**	-0.786**
	(0.363)	(0.364)	(0.348)	(0.353)	(0.353)	(0.352)
Lead×Output and reforms indicator	1.250**	1.252**	1.800**	1.778**	1.075*	1.092*
	(0.619)	(0.623)	(0.856)	(0.850)	(0.634)	(0.632)
R-squared	0.155	0.152	0.157	0.152	0.154	0.152
Observations	1187	1187	1187	1187	1187	1187
No. countries	56	56	56	56	56	56

Notes: The dependent variable is the change in the financial reform index. *Output indicator* is an indicator variable equal to 1 when lagged output per capita is greater than its lagged trend value. The trend value of output per capita is obtained using the HP filter with a smoothing parameter of 100. *Output and reforms indicator* is an indicator variable equal to 1 when *Output indicator* is 1 and the lagged change in financial reform index is greater than or equal to zero. All regressions include country and year fixed effects. Standard errors are in parentheses, clustered at the country level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level.

election is held later or earlier than its pre-determined date). It also includes cases when the constitution is suddenly changed, creating a disjuncture between the prevailing term length and the new mandated term length. It is important to consider these latter cases as deviant cases since a leader may be able to engineer a constitutional change and alter term length in response to, for example, political and economic conditions in the country. Imposing this sample restriction causes the sample size to fall substantially, by almost one-fifth. Appendix

Table B.3: Robustness for executive elections

	Pre-determined			First term			Second term or above		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Lag	1.047*	1.221**		1.052*	1.315**		0.779	0.710	
	(0.538)	(0.509)		(0.600)	(0.590)		(0.701)	(0.721)	
Year of	-0.002	0.167	-0.344	0.330	0.609	0.001	-0.661	-0.732	-0.889
	(0.439)	(0.415)	(0.466)	(0.545)	(0.531)	(0.509)	(0.817)	(0.766)	(0.768)
Lead	-0.505		-0.849*	-0.804		-1.121*	0.254		0.023
	(0.521)		(0.487)	(0.668)		(0.648)	(0.432)		(0.493)
R-squared	0.173	0.173	0.170	0.147	0.144	0.143	0.259	0.259	0.256
Observations	989	989	989	806	806	806	389	389	389
No. countries	55	55	55	55	55	55	40	40	40

Notes: The dependent variable is the change in the financial reform index. All regressions include country and year fixed effects. Standard errors are in parentheses, clustered at the country level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level.

Table B.3 columns 1-3 show that though some of the coefficients are less precisely estimated, overall the results for executive elections largely remain in place even with this restriction.

Second, in Appendix Table B.3 columns 4-6 and columns 7-9, we compare chief executives in the first term in office with latter terms in office. To do so, we gathered information on each chief executive's name in our sample, to track whether a term actually marked their first time holding office. (This additional step is important since multiple terms are not always consecutive.) Consistent with the model's implication, we observe much stronger evidence of reform fatigue during an executive's first term in office (columns 4-6), as opposed to latter terms in office (columns 7-9). Coefficients on the lag and lead indicator variables are both smaller and statistically insignificant during the latter terms. In fact, the coefficient on the lead variable is even opposite in sign (i.e., positive) in columns 7 and 9.

B.3 Reform Fatigue and the IMF

We also consider whether financial reform cycles are influenced by participation in IMF programs. We attain information about countries' historical lending arrangements with the IMF from the IMF's website (<https://www.imf.org/external/np/fin/tad/extarr1.aspx>). We create an indicator variable, *IMF*, which equals one if a country has an outstanding loan, yet to expire, from the Fund.²⁸ IMF programs are quite common in our sample, with just over half of the country-year observations falling under such a program. In Appendix Table B.4 we present estimates interacting our election variables with the indicator of IMF program

²⁸Loans can be any of the following types: Exogenous Shock Facility, Extended Credit Facility, Extended Fund Facility, Flexible Credit Line, Precautionary and Liquidity Line, Precautionary and Liquidity Line, Standby Arrangement, Standby Credit Facility, and Structural Adjustment Facility Commitment.

Table B.4: Financial reforms and IMF program participation

	Executive or Legislative			Executive			Legislative		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Lag	0.835 (0.538)	0.953* (0.541)		1.043 (0.700)	1.159* (0.698)		0.592 (0.624)	0.743 (0.623)	
Lag×IMF	-0.213 (0.697)	-0.127 (0.699)		-0.050 (0.878)	0.024 (0.861)		0.068 (0.732)	0.248 (0.748)	
Year of	0.612 (0.655)	0.764 (0.618)	0.344 (0.564)	0.255 (0.584)	0.371 (0.564)	-0.064 (0.529)	0.883 (0.752)	1.079 (0.694)	0.664 (0.620)
Year of×IMF	-0.994 (0.872)	-0.846 (0.816)	-0.943 (0.778)	-0.176 (0.873)	-0.095 (0.806)	-0.169 (0.804)	-1.531 (0.962)	-1.299 (0.896)	-1.543* (0.831)
Lead	-0.310 (0.480)		-0.485 (0.478)	-0.299 (0.609)		-0.586 (0.598)	-0.405 (0.545)		-0.545 (0.541)
Lead×IMF	-0.613 (0.658)		-0.606 (0.649)	-0.374 (0.733)		-0.405 (0.719)	-0.852 (0.717)		-0.897 (0.704)
IMF	2.223*** (0.675)	2.002*** (0.586)	2.164*** (0.568)	1.864*** (0.503)	1.787*** (0.471)	1.846*** (0.477)	2.234*** (0.581)	1.953*** (0.522)	2.283*** (0.508)
R-squared	0.173	0.170	0.170	0.171	0.170	0.167	0.176	0.171	0.174
Observations	1195	1195	1195	1195	1195	1195	1195	1195	1195
No. countries	56	56	56	56	56	56	56	56	56
<i>F</i> -test (<i>p</i>)	0.634	0.519	0.421	0.966	0.990	0.853	0.191	0.124	0.137

Notes: The dependent variable is the change in the financial reform index. All regressions include country and year fixed effects. The row *F*-test (*p*) shows p-values associated with the *F*-test for the joint significance of *Lag* × *IMF*, *Year of* × *IMF*, and *Lead* × *IMF*. Standard errors are in parentheses, clustered at the country level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level.

participation. The coefficient on *IMF* is positive and significant as expected, indicating that countries under IMF programs do in fact, implement financial reforms to a greater degree. However, the coefficients on the interaction terms with the IMF variable are insignificant across specifications. This suggests that the political cycle in reforms is neither dampened nor exacerbated by program participation. Thus, reform fatigue appears to not be driven by external pressure from this international organization.

B.4 Financial reforms in parliamentary regimes

In Appendix Table B.5, we present our main specification but for parliamentary countries. Parliamentary regimes typically do not have an executive election that is separate from the legislative election, so we present results for the combined elections and legislative elections only. We find no evidence of a political cycle in financial reforms in this sample. Since executives are not elected directly under parliamentary systems, this is consistent with our theoretical framework, which posits that reform cycles emerge when politicians are directly accountable to voters.

Table B.5: Financial reforms in parliamentary regimes

	Executive or Legislative			Legislative		
	(1)	(2)	(3)	(4)	(5)	(6)
Lag	0.415 (0.419)	0.301 (0.341)		0.602 (0.444)	0.419 (0.363)	
Year of	0.187 (0.327)	0.065 (0.267)	-0.001 (0.252)	0.338 (0.346)	0.141 (0.278)	0.056 (0.251)
Lead	0.270 (0.349)		0.093 (0.269)	0.423 (0.369)		0.160 (0.285)
R-squared	0.089	0.088	0.088	0.090	0.089	0.088
Observations	909	909	909	909	909	909
No. countries	41	41	41	41	41	41

Notes: The dependent variable is the change in the financial reform index. All regressions include country and year fixed effects. Standard errors are in parentheses, clustered at the country level. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level.

References

- Abiad, A., Detragiache, E. and Tressel, T. (2008), ‘A New Database of Financial Reforms’, *IMF Working Paper Working No. 08/266*.
- Abiad, A. and Mody, A. (2005), ‘Financial Reform: What Shakes it? What Shapes it?’, *American Economic Review* **95**(1), 66–88.
- Ales, L., Maziero, P. and Yared, P. (2014), ‘A Theory of Political and Economic Cycles’, *Journal of Economic Theory* **153**, 224–251.
- Alesina, A. and Cukierman, A. (1990), ‘The Politics of Ambiguity’, *Quarterly Journal of Economics* **105**(4), 829–850.
- Ashworth, S. (2005), ‘Reputational Dynamics and Political Careers’, *Journal of Law, Economics & Organization* **21**(5).
- Ashworth, S. (2012), ‘Electoral Accountability: Recent Theoretical and Empirical Work’, *Annual Review of Political Science* **15**, 183–201.
- Backus, D. K. and Kehoe, P. J. (1992), ‘International Evidence of the Historical Properties of Business Cycles’, *American Economic Review* **82**(4), 864–888.
- Banks, J. S. (1998), ‘Optimal Retention in Agency Problems’, *Journal of Economic Theory* **82**, 293–323.

- Barro, R. J. and Ursúa, J. F. (2008), ‘Macroeconomic Crises since 1870’, *Brookings Papers on Economic Activity* **39**(1), 255–350.
- Berliant, M. and Duggan, J. (2010), ‘Dynamic Political Agency with Adverse Selection and Moral Hazard’, *Working Paper* .
- Besley, T. and Case, A. (1995), ‘Does Electoral Accountability Affect Economic Policy Choices? Evidence from Gubernatorial Term Limits’, *Quarterly Journal of Economics* **110**(3), 769–798.
- Besley, T. and Coate, S. (1998), ‘Sources of Inefficiency in a Representative Democracy: A Dynamic Analysis’, *American Economic Review* **88**(1), 139–156.
- Bonatti, A. and Hörner, J. (Forthcoming), ‘Career Concerns with Exponential Learning’, *Theoretical Economics* .
- Bowen, T. R. and Mo, C. H. (Forthcoming), ‘The Voter’s Blunt Tool’, *Journal of Theoretical Politics* .
- Brender, A. and Drazen, A. (2005), ‘Political budget cycles in new versus established democracies’, *Journal of Monetary Economics* **52**, 1271–1295.
- Callander, S. and Harstad, B. (2015), ‘Experimentation in Federal Systems’, *Quarterly Journal of Economics* **130**(2), 951–1002.
- Canes-Wrone, B., Herron, M. C. and Shotts, K. W. (2001), ‘Leadership and Pandering: A Theory of Executive Policymaking’, *American Journal of Political Science* **45**(3), 532–550.
- Canes-Wrone, B. and Shotts, K. W. (2004), ‘The Conditional Nature of Presidential Responsiveness to Public Opinion’, *American Journal of Political Science* **48**(4), 690–706.
- Christiansen, L., Schindler, M. and Tressel, T. (2013), ‘Growth and Structural Reforms: A New Assessment’, *Journal of International Economics* **89**, 347–356.
- Cooley, T. and Ohanian, L. (1991), ‘The Cyclical Behavior of Prices’, *Journal of Monetary Economics* **28**(1), 25–60.
- Correia, I. H., Neves, J. L. and Rebelo, S. T. (1992), ‘Business Cycles from 1850-1950: New Facts about Old Data’, *European Economic Review* **36**(2-3), 459–467.
- DeMarzo, P. and Sannikov, Y. (2014), ‘Learning, Termination, and Payout Policy in Dynamic Incentive Contracts’, *Working Paper* .

- Dewatripont, M. and Roland, G. (1995), ‘The Design of Reform Packages under Uncertainty’, *American Economic Review* **85**(5), 1207–1223.
- Dewatripont, M. and Roland, G. R. (1992), ‘Economic Reform and Dynamic Political Constraints’, *Review of Economic Studies* **59**, 703–730.
- Di Tella, R. and MacCulloch, R. J. (2005), ‘The Consequences of Labor Market Flexibility: Panel Evidence Based on Survey Data’, *European Economic Review* **49**(5), 1225–59.
- Drazen, A. (2001), The Political Business Cycle after 25 Years, *in* B. S. Bernanke and K. Rogoff, eds, ‘NBER Macroeconomics Annual 2000’, Vol. 15, pp. 75–138.
- Duggan, J. and Forand, J. (2014), ‘Dynamic Elections and the Limits of Accountability’, *Working Paper* .
- Faye, M. and Niehaus, P. (2012), ‘Political Aid Cycles’, *American Economic Review* **102**(7), 3516–3530.
- Feldmann, H. (2012), ‘Product Market Regulation and Labor Market Performance around the World’, *Labour* **26**(3), 369–391.
- Fernandez, R. and Rodrik, D. (1991), ‘Resistance to Reform: Status Quo Bias in the Presence of Individual-Specific Uncertainty’, *American Economic Review* **81**(5), 1146–1155.
- Forteza, A., Tommasi, M. and Herrera, G. (2006), ‘Understanding reform in latin america’, *Global Development Network Global Research Project on Understanding Reform: Synthesis of Country Studies from Latin America* .
- Freidel, F. and Sidey, H. (2006), *The Presidents of the United States of America*, White House Historical Association.
- Garfagnini, U. (2011), ‘Delegated Experimentation’, *Working Paper* .
- Giuliano, P., Mishra, P. and Spilimbergo, A. (2013), ‘Democracy and Reforms: Evidence from a New Dataset’, *American Economic Journal: Macroeconomics* **5**(4), 179–204.
- Golder, M. (2005), ‘Democratic Electoral Systems Around the World, 1946-2000’, *Electoral Studies* **24**, 103–121.
- Guo, Y. (Forthcoming), ‘Dynamic Delegation of Experimentation’, *American Economic Review* .

- Halac, M., Kartik, N. and Liu, Q. (Forthcoming), ‘Optimal Contracts for Experimentation’, *Review of Economic Studies* .
- Hodrick, R. and Prescott, E. (1981), ‘Post-War U.S. Business Cycles: An Empirical Investigation’, *Carnegie Mellon University Discussion Paper No. 451* .
- Hodrick, R. and Prescott, E. (1997), ‘Postwar U.S. Business Cycles: An Empirical Investigation’, *Journal of Money, Credit and Banking* **29**(1), 1–16.
- Holmstrom, B. (1999), ‘Managerial Incentive Problems: A Dynamic Perspective’, *Review of Economic Studies* **66**(1), 169–182.
- Hörner, J. and Samuelson, L. (2013), ‘Incentives for Experimenting Agents’, *RAND Journal of Economics* **44**, 632–663.
- Jackson, M. and Aghion, P. (Forthcoming), ‘Inducing Leaders to Take Risky Decisions: Dismissal, Tenure and Term Limits’, *American Economic Journal: Microeconomics* .
- Keller, G., Rady, S. and Cripps, M. (2005), ‘Strategic Experimentation with Exponential Bandits’, *Econometrica* **73**(1), 39–68.
- Klein, N. (Forthcoming), ‘The Importance of Being Honest’, *Theoretical Economics* .
- Krueger, A. (1993), *Political Economy of Policy Reform in Developing Countries*, MIT Press: Cambridge, MA.
- Lizzeri, A. and Persisco, N. (2011), *Electoral Incentives, Political Risk-Taking and Policy Reform*, Fundación BBVA.
- Lora, E., Panizza, U. and Quispe-Agnoli, M. (2003), ‘Reform Fatigue: Symptoms, Reasons and Implications’, *Inter-American Development Bank, Research Department* .
- Manso, G. (2011), ‘Motivating Innovation’, *Journal of Finance* **LXVI**(5).
- Nordhaus, W. D. (1975), ‘The Political Business Cycle’, *Review of Economic Studies* **42**(2), 169–190.
- Prati, A., Onorato, M. G. and Papageorgiou, C. (2013), ‘Which Reforms Work and Under What Institutional Environment: Evidence from a New Dataset on Structural Reforms’, *Review of Economics and Statistics* **95**(3), 946–968.

- Ravn, M. O. and Uhlig, H. (2002), 'On Adjusting the Hodrick-Prescott Filter for the Frequency of Observations', *Review of Economics and Statistics* **84**(2), 371–375.
- Rogoff, K. (1990), 'Equilibrium Political Budget Cycles', *American Economic Review* **80**(1), 21–36.
- Shi, M. and Svensson, J. (2006), 'Political Budget Cycles: Do They Differ Across Countries and Why?', *Journal of Public Economics* **90**, 1367–1389.
- Strulovici, B. (2010), 'Learning While Voting: Determinants of Collective Experimentation', *Econometrica* **78**(3), 933–971.
- Tornell, A. (1998), 'Reform from Within', *NBER Working Paper No. 6497* .