

Abstract

Reputational Concerns in Political Agency Models

Andrew Yuichi Lemon

2005

Citizens in democracies can be thought of as hiring politicians to run the government. In this sense, elections are contracts. The citizens are the principals and retain politicians as agents who are responsible for making policy decisions. Provided politicians value holding office in the future, elections generate incentives for politicians to appear honest and able, and deter politicians from choosing corrupt policies. This literature is known as political agency (e.g., Ferejohn 1986)—it applies principal-agent models to analyze the behavior of politicians and the efficiency of their policies when citizens face adverse-selection and moral-hazard problems in their election decisions.

Chapters 1 and 2 of this dissertation further develop political-agency models to include strategic behavior by the challenger: unlike other political-agency models where the challenger is a static alternative to the incumbent, the challenger acts to maximize her payoffs. The first chapter asks whether political debate can improve the efficiency of government policies. The second chapter builds a Bayesian updating model to examine the incentives behind policy flip-flopping. The third chapter formalizes a political economy critique against strategic trade policy using a political-agency model.

Reputational Concerns in Political Agency Models

A Dissertation
Presented to the Faculty of the Graduate School
of
Yale University
in Candidacy for the Degree of
Doctor of Philosophy

by
Andrew Yuichi Lemon

Dissertation Director: Stephen Morris

May 2005

Copyright © 2005 by Andrew Yuichi Lemon

All rights reserved.

Contents

1	Political Debate: A Device For Disciplining Politicians	1
1.1	Introduction	3
1.2	The Model	10
1.3	The Effect of Political Debate	15
1.3.1	Incumbent as Agenda-Setter	17
1.3.2	No Agenda-Setter	17
1.3.3	Challenger as Agenda-Setter	23
1.4	Conclusion	35
2	Flip-floppers versus Hard-liners	53
2.1	Introduction	55
2.2	The Model	56
2.3	Policy Positioning	61
2.4	Conclusion	78
3	Strategic Export Promotion in the Presence of Special Interests	81
3.1	Introduction	82
3.2	The Model	84
3.3	Equilibrium Subsidies and Direct Transfers	91
3.3.1	Separating Equilibria	95

3.3.2	Partial-Pooling Equilibria	100
3.3.3	Discussion	102
3.4	A Continuum of Cournot Industries	104
3.5	Conclusion	108

List of Figures

1	No Agenda-Setter	22
2	Challenger as Agenda-Setter (Case 1)	35
3	Challenger as Agenda-Setter (Case 2)	51
4	Challenger as Agenda-Setter (Case 3)	51
5	Challenger as Agenda-Setter (Case 4)	52
6	Challenger as Agenda-Setter	52
1	Time Line	60
2	Politician Posteriors	63
3	Opportunistic Policy Flop-flopping	71
4	Opportunistic Policy Maintenance	74
5	Two Numerical Examples	77

Acknowledgements

I am especially grateful to my advisors, Stephen Morris and Benjamin Polak, for their guidance and encouragement throughout this project. I also wish to thank Stephen Coate, Dino Gerardi, Andrew Postlewaite, Juan Moreno-Ternero, Phil Levy, and the participants of the Game Theory Lunch for their helpful comments.

Chapter 1

Political Debate:

A Device For Disciplining

Politicians

Keywords: political debate, political agency, corruption, elections, cheap-talk, signaling

JEL classification: D72, D73, D82

Abstract

Is an elected official's policy decision corrupt? This question frequently confronts citizens of democracies, particularly when they are less informed than politicians regarding policy issues. Under such an informational asymmetry, politicians can act venally and citizens will be none the wiser. We show that when politicians debate policies, the citizen can learn which policies are corrupt and use this information to screen for corrupt politicians and discipline their policy choices. While there is always a possibility that debate is uninformative and payoff-irrelevant, we identify environments where debate can improve the citizens' welfare and where it can surprisingly harm the citizens' welfare. The effects depend crucially on: (i) which politician is the agenda-setter, i.e., who stakes out a policy position first, (ii) the benefit politicians receive from implementing their preferred policies relative to winning election, and (iii) each politician's reputation.

1.1 Introduction

From the developing world to the developed world, corruption among elected officials in local or national government is a major concern. Citizens are often less informed on policy issues than their elected officials and, therefore, cannot be fully confident that elected officials make decisions that are free of venality. In representative democracies, there are institutions which serve to detect and prevent corrupt political behavior—one such institution is political debate. When politicians debate policy issues before the public, citizens become more informed about issues and can learn which policies are corrupt. Thus, elected officials who choose corrupt policies can be identified and punished. To avoid being punished, corrupt politicians may instead choose policies which benefit citizens. This optimistic view suggests that political debate helps citizens to recognize corrupt politicians and constrain their behavior.

However, there is no guarantee that political debate will succeed in its watchdog role. The problem is that politicians who provide information in political debates are not disinterested parties, but have their own objectives. They may themselves be corrupt and use debates to further their own policy objectives, or they may be overly concerned with winning elections and use debates to curry favor with citizens. In general, it is not clear that politicians will sincerely inform citizens. This cynical view suggests that political debate may be useless in identifying corrupt politicians and constraining their behavior.¹

We model political debates as situations where an incumbent politician must make a costly policy decision while a challenger makes a cheap-talk announcement staking out a position on the decision. At the local level, consider the example of a city faced with the decision of building a road: the incumbent (i.e., the mayor) must decide

¹See Mackie (1998) for a discussion on the informative role of political communication and whether it benefits citizens.

whether to build the road, while a challenger (i.e., a political rival) states whether she thinks the road should be built. At the national level, consider the example of whether tax cuts should be used to stimulate the economy: the incumbent (i.e., the president) must decide whether to cut taxes, while a challenger (i.e., a political rival) states whether she thinks taxes should be cut. We assume that the citizen has no information regarding which policy is beneficial, but that both candidates are fully informed about each policy's benefit. Therefore, debate offers an opportunity for the candidates to inform the citizen of the appropriate and corrupt policies—the citizen can learn which policy is beneficial and which is harmful. The citizen, however, questions each candidate's motives. In particular, he is uncertain about whether a politician favors corrupt policies or appropriate policies, but knows that each politician values holding office. After the incumbent's policy decision and challenger's announcement are made, but before the outcome of the policy is known, an election is held between the two candidates in which the citizen alone determines the winner by voting for the candidate he believes is more likely to favor appropriate policies.

Despite the cheap-talk nature of the challenger's announcement, we find that political debate can be informative and have some effect on policy. If the citizen learns which policy is corrupt from the debate, an incumbent who chooses the corrupt policy can be identified by the citizen and voted out of office. That is, citizens can use the debate to screen politicians. To avoid being identified, incumbents who prefer corrupt policies may choose the appropriate policy. That is, political debate can discipline the corrupt incumbent into choosing the appropriate policy. Whether the citizen learns which policy is corrupt from the debate depends on the challenger's cheap-talk announcement and there is no guarantee that the challenger will truthfully inform the citizen. She will only reveal the corrupt policy if doing so is in her interests, i.e., it must be incentive compatible.

As we show in this paper, the effect of political debate on identifying politicians and disciplining their policy choices depends on which candidate is the agenda-setter—that is, which politician stakes out a policy position first. We consider the following three cases: (i) the incumbent as the agenda-setter, (ii) no agenda-setter, and (iii) the challenger as the agenda-setter. When the incumbent is the agenda-setter (i.e., incumbent moves first), political debate never informs the citizen of the corrupt policy. Hence, political debate has no identifying effect and no disciplining effect. This result captures the cynical view that political debate is meaningless. When neither the incumbent nor challenger is the agenda-setter (i.e., incumbent and challenger move simultaneously), political debate can inform the citizen of the corrupt policy under some circumstances. If the corrupt incumbent places sufficiently high value on choosing corrupt policies, she will choose the corrupt policy and be identified by the citizen. If instead the corrupt incumbent places sufficiently high value on holding office, she will choose the appropriate policy with some positive probability. Here political debate allows the citizen to screen incumbents and discipline their policy decisions. This result captures the optimistic view that political debate is beneficial for society. When the challenger is the agenda-setter (i.e., challenger moves first), we again find similar effects to those when there is no agenda-setter, but we also find effects which harm the citizen. These detrimental effects arise from a corrupt challenger’s ability to set the debate and an incumbent’s willingness to sacrifice policy for better prospects of winning the election. This result illustrates the dangers inherent in the ability politicians have to use political debates to further their own objectives.

Our model of political competition is related to the literature on political agency. The citizen is the principal and the incumbent politician is the agent who must make policy decisions under pressure of reelection. In early models of policy agency, citizens discipline incumbent politicians into choosing appropriate policies

by rewarding them with reelection if their performance in office is satisfactory—hence voting behavior is *retrospective*.² Ferejohn’s (1986) paper is a classic work in this spirit. He considers a setting of repeated elections in which all politicians are identical and models the political agency problem as one of moral hazard. There is a single citizen whose payoff is a random, but increasing function of the incumbent’s effort. The incumbent benefits from holding office, but finds higher levels of effort more costly. The citizen can only observe a noisy signal of the incumbent’s effort. Ferejohn shows that the citizen can induce the incumbent into choosing greater levels of effort by using a cutoff voting rule in which the incumbent is reelected if and only if the signal of the incumbent’s effort satisfies some threshold.

More recent models of political agency have loosened the unrealistic assumption that politicians are identical and instead consider political agency with a variety of politician types. When politicians are heterogenous, citizens can no longer commit to retrospective voting rules. Regardless of a politician’s past behavior, at the election citizens have an incentive to vote for the politician they believe will give them the greatest future payoff—hence voting behavior is *prospective* and incumbents are disciplined through reputational effects. Banks and Sundaram (1998) examine just such a political agency problem in a model where a single citizen elects a politician who can hold office for at most two periods. There is a moral hazard problem akin to that in Ferejohn’s paper, but there is also an adverse selection problem as the citizen is uncertain about each politician’s type. They show that there exists a sequentially rational cutoff voting rule in which the incumbent is reelected with probability one if the signal of the incumbent’s action satisfies some threshold. The cutoff rule acts as a screening mechanism—the citizen believes with probability one that the incumbent is of a desirable type if and only if the noisy signal is above the threshold. In

²Barro’s (1973) paper is among the earliest to consider the threat of reelection as a means for controlling politicians.

addition, Banks and Sundaram show that the incumbent's effort level is higher in her first period in office than in her second—that is, the citizen's forward-looking voting rule disciplines the first-period incumbent into choosing the citizen's more preferred effort level.

This literature has left the role of the challenger unexplored—one contribution of this paper is to develop the role of the challenger in a political agency model with moral hazard and adverse selection in which the citizen can infer nothing about the incumbent's type from her policy choice. In our model, the challenger engages the incumbent in a political debate the result of which can provide the citizen with a way to use the incumbent's policy choice as a screening device that can discipline politicians. Unlike other models of political agency, the challenger is not an inanimate alternative to the incumbent, but a strategic player seeking to maximize her own objectives.

The screening and disciplining effects in our model are related to those in Maskin and Tirole (2004) and Morris (2000). In the context of constitutional design, Maskin and Tirole show how the citizen can use the incumbent's policy choice to screen politicians and how this can lead politicians to pander to public opinion. Here the screening and disciplining effects arise solely from costly actions. In the context of cheap-talk policy advice, Morris shows how the attempts of an uninformed policymaker to screen the policy preferences of her partially informed advisor can lead biased advisors into giving truthful advice and also lead unbiased advisors into giving false advice. Here the screening and disciplining effects arise solely from cheap-talk messages. In contrast, these effects arise in our model from the interaction of costly actions and cheap-talk messages.

Our model is related to a view of political debate put forth by Austen-Smith (1990) where debate is a deliberative process in which policymakers share information

in order to improve the quality of the eventual policy. Austen-Smith investigates a direct democracy setting in which each member of a committee must propose a policy and subsequently vote over the set of proposed policies. Each member has private but correlated information regarding the benefit of various policies. Prior to proposing policies, members take part in debate over policies in which each member sends a cheap-talk message regarding her private information. Austen-Smith examines if information is revealed in equilibrium and if so whether it is revealed in the debating stage or the proposal stage of the game. Our model differs in that we investigate the effect of political debate in a representative democracy where citizens elect politicians to policymaking positions. We do not examine debate between citizens on committees, but rather debate between politicians and citizens in the public sphere.

Our modeling of cheap-talk political messages and their interaction with policy decisions is related to the study of campaign rhetoric by Aragonés, Palfrey, and Postlewaite (2004). They investigate cheap-talk campaign promises in a setting where politicians have Downsian-style ideological policy preferences and face repeated elections. In their model, two politicians competing for office make cheap-talk promises regarding the policies they will enact once elected—politicians cannot commit to the policy they will choose if elected. They show that there exist equilibria in which politicians always honor their campaign promises provided their promised policies are not too far from their most preferred policies. Moreover, they show that such equilibria yield the median voter greater utility than if all voters ignore all cheap-talk campaign promises. Our model differs in that we examine the effect of cheap-talk political messages on government policy made prior to an election.

Our paper also contributes to the literature on deliberative democracy which has received much attention in political science. Broadly speaking, this literature seeks to understand the role that communication plays in making collective decisions

through a democratic political process. Habermas (1996), Bohman (1996), and Gutmann and Thompson (1996) address the political theory question concerning how a deliberative process gives legitimacy to political decisions. Elster (1998) examines the factors that encourage the creation of a deliberative setting. Fearon (1998) discusses several reasons why a group of individuals would want to discuss matters before making a collective decision. Stokes (1998) provides examples of situations where public deliberation has led to perverse outcomes and suggests rules that would help guard against these dangers. Much of the deliberative democracy literature, including all of the works cited above, does not involve formal modeling. An exception is Stasavage (2004) which considers the trade-offs between public and private deliberation in a representative democracy.

Our model of political debate is related to both Fearon (1998) and Stokes (1998). Fearon states that discussion is desirable because it gives the participants an opportunity to reveal private information. Although Fearon emphasizes the sharing of private information among citizens in a direct democracy or among representatives in an assembly, we show how politicians can credibly communicate private information to citizens in the context of political debates within representative democracies. As mentioned above, we find that there are circumstances in which political debate leads to the adoption of corrupt policies—outcomes that are undesirable for the citizen. The existence of these outcomes provides further support for the adverse effects raised by Stokes (1998).

This paper is organized as follows. In section 2, we present our model. In section 3, we analyze the equilibria of the model in each of the agenda-setting scenarios mentioned above and discuss the impact of debate on the citizen's welfare. We illustrate the equilibria of our model for the special case where the citizen's voting rule is linear in his beliefs regarding the politicians' motivations. This allows us to

show graphically where certain equilibria exist in the parameter space. Section 4 concludes.

1.2 The Model

We construct a model in which an incumbent politician is seeking reelection against a challenger. Prior to the election, the incumbent must make a policy decision and the challenger must make a costless policy recommendation. The citizen does not know which policy will benefit him and does not know which politician, if any, shares his policy preferences.³ Each politician knows which policy benefits the citizen, but shares the citizen’s uncertainty regarding the other politician’s policy preferences. After the policy and policy suggestion have been made, an election is held between the incumbent and challenger. The citizen alone determines the winner.

Government Policies and Policy Uncertainty. There are two policies, $\underline{\theta}$ and $\bar{\theta}$. One of these policies benefits the citizen—we call this policy “appropriate” and label it T . The other policy does not benefit the citizen—we call this policy “inappropriate” and label it F . At the start of the game, Nature determines which policy is appropriate. The politicians observe Nature’s choice while the citizen only knows the prior that each policy is equally likely to be appropriate. The citizen is informed about the appropriateness of the policies only after the election. In this manner, the citizen faces policy uncertainty.

Politicians and Politician Uncertainty. As mentioned above, the incumbent must choose a policy $a \in \{T, F\}$ and the challenger must choose a policy suggestion $m \in \{T, F\}$.⁴ The incumbent’s policy choice is directly payoff relevant. The chal-

³Coate and Morris (1996) present a model of political agency in which the citizen faces both policy uncertainty and politician uncertainty in order to study the form of transfers to special interests.

⁴We restrict ourselves to studying politician strategies that are symmetric across states. The

lenger's policy suggestion is not directly payoff relevant, i.e., it is cheap talk. We will consider three strategic situations between the incumbent and challenger: (i) the incumbent as the agenda-setter where the incumbent's policy decision is made before the challenger's policy recommendation, (ii) the challenger as the agenda-setter where the challenger's policy recommendation is made before the incumbent's policy decision, and (iii) no agenda-setter where the challenger's policy recommendation is made simultaneously to the incumbent's policy decision.⁵

Whether a politician is an incumbent or a challenger, her payoffs depend on her type: *good* or *bad*. Good politicians have the same policy preferences as the citizen; that is, the appropriate policy benefits good politicians while the inappropriate policy does not. Specifically, a good politician receives payoff $b_G > 0$ if the appropriate policy is chosen by the incumbent. She receives payoff 0 if the inappropriate policy is chosen by the incumbent. Bad politicians have the opposite preferences; that is, the inappropriate policy benefits bad politicians while the appropriate policy does not. Specifically, a bad politician receives payoff $b_B > 0$ if the inappropriate policy is chosen by the incumbent. She receives payoff 0 if the appropriate policy is chosen by the incumbent.

In addition to their policy benefits, both the incumbent and the challenger want to hold office, regardless of their type. If a politician is elected/re-elected she receives an additional payoff which we normalize to 1; if she is not elected/re-elected, she receives no additional payoff. This normalization allows us to interpret b_G and b_B , respectively, as the value the good and bad politician place on their preferred policy

politicians could employ strategies that dictate different behavior when policy $\underline{\theta}$ is appropriate and when policy $\bar{\theta}$ is appropriate. Such strategies would complicate analysis.

⁵These three cases can be thought of reflecting which politician's policy advisory group discovers an important policy issue first. Some times the incumbent's group discovers the important policy issue first. Other times the challenger's group discovers the important policy issue first. And on other occasions, both groups discover the important policy issue at the same time when it suddenly appears in the headlines of the newspapers.

today relative to holding office tomorrow—the greater is b_G (or b_B), the more the good (or bad) politician values her preferred policy relative to holding office tomorrow.

Each politician’s payoffs can be concisely represented by her utility function. The good incumbent’s utility function $U_G(a; m)$ and the bad incumbent’s utility function $U_B(a; m)$ are:

$$U_G(a; m) = \begin{cases} b_G + \Pr(\text{reelect} | a = T, m) & \text{if } a = T \\ 0 + \Pr(\text{reelect} | a = F, m) & \text{if } a = F \end{cases} \quad (1)$$

$$U_B(a; m) = \begin{cases} 0 + \Pr(\text{reelect} | a = T, m) & \text{if } a = T \\ b_B + \Pr(\text{reelect} | a = F, m) & \text{if } a = F \end{cases} \quad (2)$$

where $\Pr(\text{reelect} | a, m)$ refers to the probability that the incumbent is reelected conditional her policy choice a and the incumbent’s policy recommendation m . The good challenger’s utility function $V_G(m; a)$ and the bad challenger’s utility function $V_B(m; a)$ are:

$$V_G(m; a) = \begin{cases} b_G + (1 - \Pr(\text{reelect} | a = T, m)) & \text{if } a = T \\ 0 + (1 - \Pr(\text{reelect} | a = F, m)) & \text{if } a = F \end{cases} \quad (3)$$

$$V_B(m; a) = \begin{cases} 0 + (1 - \Pr(\text{reelect} | a = T, m)) & \text{if } a = T \\ b_B + (1 - \Pr(\text{reelect} | a = F, m)) & \text{if } a = F \end{cases} \quad (4)$$

Notice, that the challenger’s cheap-talk policy suggestion affects each politicians payoffs indirectly through her election/reelection probability.

The Citizen. The citizen wants to elect/reelect good politicians. However, he is uncertain about each politician’s type. Instead, he has an initial belief β_I that the incumbent is good and an initial belief β_C that the challenger is good. We respectively call these beliefs the incumbent’s and challenger’s “initial reputation.”

Prior to the election, the citizen observes the challenger's policy suggestion m and the incumbent's policy a . The citizen has no other information which he can use to update his beliefs about the politicians. As there are only two possible policies and policy recommendations, the citizen may observe one of two events: (i) an identical policy and policy recommendation which we label x , or (ii) a different policy and policy recommendation which we label y . After seeing event x , the citizen Bayesian updates his beliefs regarding the incumbent and challenger to $\widehat{\beta}_I(x)$ and $\widehat{\beta}_C(x)$. Similarly, after seeing event y , he Bayesian updates his beliefs regarding the incumbent and challenger to $\widehat{\beta}_I(y)$ and $\widehat{\beta}_C(y)$.

When the election is held, the citizen votes probabilistically for the incumbent according to the following cumulative distribution function:

$$H(\widehat{\beta}_I(\cdot), \widehat{\beta}_C(\cdot)) \tag{5}$$

which is increasing in the first term and decreasing in the second term. That is, the citizen is more likely to vote for the incumbent as the incumbent's reputation increases and is less likely to vote for the incumbent as the challenger's reputation increases.⁶

The Game and the Definition of Equilibrium. This model defines a game between the incumbent, challenger, and citizen. At the beginning of the game, Nature chooses the incumbent's type and the challenger's type. Then Nature determines which policy is appropriate. As mentioned above, the politicians observe which policy is appropriate. The citizen knows only the prior that each policy is equally likely to be appropriate. Furthermore, the citizen is uncertain about each politician's type.

⁶This probabilistic voting function can be thought of as the reduced form of a model where the citizen's voting depends on other issues, but where the politicians' are uncertain about the citizen's preferences regarding these other issues. Hence, the citizen's voting behavior is probabilistic from the politicians' point of view.

The incumbent and challenger are similarly uninformed about the other's type. That is, each politician holds the same belief as the citizen regarding the other's type. After Nature has moved, the incumbent and challenger make their policy choice and policy recommendation. The citizen observes the incumbent's policy choice and the challenger's policy recommendation, and updates his beliefs. The election is then held between the incumbent and challenger and the winner takes office.

As mentioned above, the timing of the incumbent's and challenger's moves depends on which politician, if any, is the agenda-setter. In the analysis below, we will consider the following three cases in order: (i) the incumbent as the agenda-setter, (ii) no agenda-setter, and (iii) the challenger as the agenda-setter.

In this model, a Perfect Bayesian Nash Equilibrium (PBNE) consists of: (1) a strategy for the incumbent, (2) a strategy for the challenger, (3) a strategy for the citizen, and (4) the citizen's beliefs. First, each politician's strategy must be optimal given her opponent's strategy and the beliefs and strategy of the citizen. Second, the citizen's strategy must be optimal given his beliefs. Third, the citizen's beliefs must be consistent with each politician's strategy, i.e., they must be formed according to Bayes's Rule where possible.

A strategy for the incumbent depends on the timing of the incumbent's and challenger's moves. If the incumbent moves first or moves simultaneously with the challenger, the incumbent's strategy is simply a rule which states which policy, T or F , to choose for each type the incumbent might be. If the incumbent moves after the challenger, the incumbent's strategy is a rule which states which policy, T or F , to choose after each possible policy suggestion from the challenger, $m = T$ and $m = F$, for each type the incumbent might be.

A strategy for the challenger similarly depends on the timing of the incumbent's and challenger's moves. If the challenger moves first or moves simultaneously

with the incumbent, the challenger's strategy is simply a rule which states which policy suggestion, T or F , to choose for each type the challenger might be. If the challenger moves after the incumbent, the challenger's strategy is a rule which states which policy suggestion, T or F , to choose after each policy the incumbent could choose, $a = T$ and $a = F$, for each type the challenger might be.

A strategy for the citizen is a rule that specifies the probability he elects the incumbent. This rule depends on each candidate's strategy and initial reputation. In addition the citizen's strategy, we must specify his beliefs regarding each candidate's type which also depends on each candidate's strategy and initial reputation.

1.3 The Effect of Political Debate

In all models of cheap talk, there is a babbling equilibrium in which the cheap-talk message is ignored. Regardless of the timing between the incumbent and challenger, our model is no different. Suppose good and bad challengers always randomize equally between policy recommendations T and F . As the policy recommendation is independent of which policy is appropriate, the citizen ignores the recommendation. Given that the citizen ignores the challenger's recommendation, there is no way for the citizen to update the initial reputations of the incumbent and challenger because each policy is equally likely to be appropriate. Regardless of the incumbent's policy choice and challenger's policy recommendation, the citizen continues to believe that the incumbent is good with probability β_I and that the challenger is good with probability β_C . Moreover, provided the citizen ignores the challenger's message, the challenger has no incentive to send an informative message as she cannot affect her initial reputation or that of the incumbent and hence her utility. Therefore, sending an uninformative message is a best response for the challenger. She babbles.

Now, suppose the challenger babbles. From the above discussion, we know that the citizen elects the incumbent with probability $H(\beta_I, \beta_C)$. As each incumbent's policy choice does not affect her reelection probability, each incumbent has a strict incentive to choose her preferred policy. Hence, the good incumbent chooses the appropriate policy T and bad incumbent chooses the inappropriate policy F . Our first proposition summarizes the babbling equilibrium.

Proposition 1 *In the babbling equilibrium, the good incumbent chooses the appropriate policy T and the bad incumbent chooses the inappropriate policy F . In other words, the appropriate policy is implemented with frequency β_I and the inappropriate policy with frequency $1 - \beta_I$.*

As all cheap-talk games have babbling equilibria, the interesting question is when do informative equilibria exist and how do they differ from the babbling equilibrium. We make the following distinction among the characteristics of informative equilibria: (i) politician revealing, (ii) bad-incumbent disciplining, and (iii) good-incumbent disciplining.⁷ An informative equilibrium is *politician revealing* when the challenger's message allows the citizen to identify good and bad politicians. An informative equilibrium is *bad-incumbent disciplining* when the challenger's message has the affect of increasing the frequency the bad incumbent chooses the appropriate policy vis-à-vis the babbling equilibrium. An informative equilibrium is *good-incumbent disciplining* when the challenger's message has the affect of increasing the frequency the good incumbent chooses the inappropriate policy vis-à-vis the babbling equilibrium.

We turn now to the task of finding informative equilibria. First, we examine the game in which the incumbent is the agenda-setter. Second, we investigate the

⁷Morris (2000) refers to similar characteristics, respectively, as the sorting effect, the disciplining effect, and the political correctness effect.

game in which neither candidate is the agenda-setter. Third, we analyze the game in which the challenger is the agenda-setter.

1.3.1 Incumbent as Agenda-Setter

Suppose the incumbent makes her policy choice before the challenger makes her policy recommendation. Once the incumbent has made her policy choice, the challenger's message alone determines which event, x or y , the citizen observes. The challenger has a strict incentive to choose the policy recommendation which leads to the event giving her a higher election probability and hence a greater expected election payoff. Therefore, the challenger's policy suggestion is always independent of which policy is appropriate. Her message is always uninformative—she babbles. In this setting, there are no informative equilibria. Political debate has no effect on the citizen's beliefs and therefore no disciplining effect on incumbent politicians. Proposition 2 summarizes this discussion.

Proposition 2 *When the incumbent's policy choice precedes the challenger's policy suggestion, the challenger always babbles. Therefore, there are no informative equilibria when the incumbent is the agenda-setter.*

1.3.2 No Agenda-Setter

Suppose the incumbent makes her policy choice at the same time the challenger makes her policy recommendation. In this setting, the good and bad challengers are only able to influence the citizen's beliefs, not the incumbent's policy choice. Therefore, the good and bad challengers consider only their expected election payoffs when deciding between message T and F . Given the initial reputations β_I and β_C , good and bad challengers have the same election probabilities. This means that good

and bad challengers face the same incentives when deciding which message to send. Therefore, they pool and effectively there is only one challenger. For this reason, we do not make a distinction between the good and bad challengers in the following analysis and instead simply refer to a single challenger.

When the incumbent makes her policy choice at the same time as the challenger makes her policy recommendation, we find a generic informative equilibrium that is incumbent revealing given certain parameters and another generic informative equilibrium that is partially incumbent revealing and bad-incumbent disciplining given certain other parameters. For all remaining parameters, only the babbling equilibrium exists. These equilibria are described in Proposition 3.

Proposition 3 *Depending on the parameters, one of two generic informative equilibria may exist.*

1. **Incumbent Revealing.** *There is an informative equilibrium in which the challenger suggests T (w.l.o.g.⁸), the good incumbent implements T , and the bad incumbent implements F , provided:*

$$b_B \geq H(1, \beta_C) - H(0, \beta_C) \tag{3.1-B}$$

$$\frac{1}{2} \geq \beta_I \tag{3.1-C}$$

2. **Partially Incumbent Revealing and Bad-Incumbent Disciplining.** *There is an informative equilibrium in which the challenger suggests T (w.l.o.g.), the good incumbent implements T , and the bad incumbent mixes over the policies*

⁸There are informative equilibria with identical incumbent behavior in which the challenger recommends policy F . In these equilibria, the following analysis goes through with the change that the beliefs attached to events x and y are reversed. The citizen still knows which policy is appropriate—the policy the challenger *does not* recommended.

with probability $0 < q^* < 1$ on T , provided:

$$H(1, \beta_C) - H(0, \beta_C) > b_B > H(\beta_I, \beta_C) - H(0, \beta_C) \quad (3.2-B)$$

$$\frac{1}{2} \left(1 - \frac{q^*}{1 - q^*} \right) \geq \beta_I \quad (3.2-C)$$

where q^* is implicitly given by

$$b_B = H(\widehat{\beta}_I(x|q^*), \beta_C) - H(0, \beta_C)$$

3. For all other parameter values, only the babbling equilibrium exists.

In the first equilibrium, Proposition 3.1, the challenger's policy recommendation reveals the appropriate policy to the citizen, but does not affect the good and bad incumbents' policy choices vis-à-vis the babbling equilibrium. Although the incumbents' behavior does not change, the citizen is able to identify the good and bad incumbents because the good and bad incumbents separate—the good incumbent chooses the appropriate policy and the bad incumbent chooses the inappropriate policy. If the citizen observes event x , he believes the incumbent is good with probability one and votes for the incumbent with probability $H(1, \beta_C)$. If he observes event y , he believes the incumbent is good with probability zero and votes for the incumbent with probability $H(0, \beta_C)$.

The behavior of each type of incumbent in Proposition 3.1 is determined by the trade-offs she faces between policy and winning election. The good incumbent chooses the appropriate policy because it is her strict best response to the challenger's informative message—she strictly benefits from the gain of her policy payoff and increase in her expected reelection payoff. The bad incumbent chooses the inappropriate policy because her incentive compatibility constraint (3.1-B) states that it is her best

response to the challenger’s informative message—she benefits more from the gain of her policy payoff than the loss she suffers from the decrease in her expected reelection payoff. The challenger will send an informative message if doing so gets her elected with greater probability than when she babbles. Notice that the probability the challenger is elected when she faces a bad incumbent in this informative equilibrium, $1 - H(0, \beta_C)$, is greater than the probability she is elected in the babbling equilibrium, $1 - H(\beta_I, \beta_C)$. Hence, sending an informative message satisfies the challenger’s incentive compatibility constraint (3.1-C) when the challenger believes that the incumbent is more likely to be bad, i.e., when the incumbent’s reputation is below $1/2$.

In the second equilibrium, Proposition 3.2, the challenger’s policy recommendation reveals the appropriate policy to the citizen and affects the bad incumbent’s policy choice vis-à-vis the babbling equilibrium. On occasion, the citizen is able to identify the bad incumbent because the bad incumbent semi-separates from the good incumbent—the good incumbent chooses the appropriate policy and the bad incumbent strictly mixes over the appropriate and inappropriate policies. If the citizen observes event x , he believes the incumbent is good with probability $\frac{\beta_I}{\beta_I + (1 - \beta_I)q^*}$ and votes for the incumbent with probability $H(\frac{\beta_I}{\beta_I + (1 - \beta_I)q^*}, \beta_C)$. If he observes event y , he believes the incumbent is good with probability zero and votes for the incumbent with probability $H(0, \beta_C)$.

Just as before, the behavior of each type of incumbent in Proposition 3.2 is determined by the trade-offs she faces between policy and winning election. The good incumbent chooses the appropriate policy because it is her strict best response to the challenger’s informative message—she strictly benefits from the gain of her policy payoff and increase in her expected reelection payoff. The bad incumbent strictly mixes over the two policies because her incentive compatibility constraint (3.2-B)

states that it is her best response to the challenger’s informative message—she is willing to forgo her policy payoff for an increase in her expected reelection payoff by choosing the appropriate policy, but as she increasingly plays the appropriate policy, her expected reelection payoff begins to fall so much so that she does not want to completely pool with the good incumbent. The challenger will send an informative message if doing so gets her elected with greater probability than when she babbles. Notice, the probability the challenger is elected when she faces a bad incumbent in this informative equilibrium, $1 - q^*(H(\widehat{\beta}_I(x|q^*), \beta_C) - H(0, \beta_C)) - H(0, \beta_C)$, is greater than the probability she is elected in the babbling equilibrium, $1 - H(\beta_I, \beta_C)$ provided the incumbent is sufficiently likely to be bad. Hence, sending an informative message satisfies the challenger’s incentive compatibility constraint (3.2-C) when the incumbent’s reputation is sufficiently low.

Proposition 3 describes the effect political debate can have when there is no agenda-setter. In particular, the first equilibrium describes the effect of political debate when bad politicians care more for inappropriate policies than they do for holding office. Here, political debate can inform the citizen of the appropriate policy if the incumbent is, on average, believed to be a bad politician. That is, debate has the potential to help the citizen exactly when the citizen is more likely to receive the inappropriate policy from the incumbent were there no debate. However, given that bad politicians favor inappropriate policies over election concerns, political debate cannot discipline bad incumbents into choosing the appropriate policy. The second equilibrium describes the effect of political debate when bad politicians value winning election, but also moderately care about inappropriate policies. Here, political debate can inform the citizen of the appropriate policy if the incumbent is sufficiently believed to be a bad politician. That is, debate has the potential to help the citizen exactly when the citizen is very likely to receive the inappropriate policy from the




Figure 1: No Agenda-Setter ($\beta < 1/2$)

incumbent were there no debate. Given that bad politicians value winning election, but also moderately care about inappropriate policies, political debate can discipline bad politicians into choosing appropriate policies. However, if bad politicians value winning election too much, political debate cannot inform the citizen of the appropriate policy and there is no disciplining effect. Hence, when there is no agenda-setter, political debate weakly improves the citizen's welfare.

Figure 1 depicts the equilibria in Proposition 3 when we assume the citizen's voting function has the following form:

$$H(\widehat{\beta}_I, \widehat{\beta}_C) = \frac{1 + \widehat{\beta}_I - \widehat{\beta}_C}{2} \quad (6)$$

Notice that there is (w.l.o.g.) at most one informative equilibrium.

At this point the reader may wonder if there exist any other generic informative equilibria for other parameter values. The following two lemmas show otherwise and

establish Proposition 3.3. The proofs are provided in the Appendix.

Lemma 1 *In any informative equilibrium, good and bad incumbents do not pool.*

Lemma 2 *In any informative equilibrium, the good incumbent does not place any positive weight on policy F .*

Together Lemma 1 and Lemma 2 show that strategy profiles other than the two considered in Proposition 3 are inconsistent with the PBNE solution concept. The intuition is that other strategy profiles imply election/reelection probabilities for which the challenger or one of the incumbents has a strict incentive to deviate. For example, consider the strategy profile in which the challenger recommends T (w.l.o.g.) and the good and bad incumbents implement T . On the equilibrium path, the citizen observes event x . The citizen's updated beliefs equal his initial beliefs and hence he votes for the incumbent with probability $H(\beta_I, \beta_C)$. Off the equilibrium path, the citizen observes event y . The citizen's updated beliefs about the politicians, $\widehat{\beta}_I(y)$ and $\widehat{\beta}_C(y)$, cannot be determined by Bayes's Rule and hence his voting rule, $H(\widehat{\beta}_I(y), \widehat{\beta}_C(y))$, cannot be tied down. The bad incumbent has no incentive to deviate from policy T if $H(\beta_I, \beta_C) > H(\widehat{\beta}_I(y), \widehat{\beta}_C(y))$. The challenger has no incentive to deviate from an informative message if $H(\widehat{\beta}_I(y), \widehat{\beta}_C(y)) \geq H(\beta_I, \beta_C)$. This is a contradiction. For any value of $H(\widehat{\beta}_I(y), \widehat{\beta}_C(y))$, the bad incumbent or the challenger has a strict incentive to deviate from the above strategy profile.

1.3.3 Challenger as Agenda-Setter

Suppose the challenger makes her policy recommendation before the incumbent makes her policy choice. Depending on the parameters, there can be one of several generic informative equilibria. We divide these equilibria into three classes: (i) equilibria that involve the bad challenger pooling with the good challenger, (ii) equilibria that

involve the bad challenger separating from the good challenger, and (iii) equilibria that involve the bad challenger partially-pooling with good challenger. The challenger pooling equilibria are outcome equivalent to the informative equilibria in the no-agenda-setter game discussed in Proposition 3 with the exception that an additional incentive compatibility constraint is required to tie down incumbent behavior. The interested reader is directed to the Appendix for analysis of this class of equilibria. The equilibria involving partial-pooling and separation of challengers have no analog in the no-agenda-setter game. Because the challenger partial-pooling equilibria are very similar to the challenger separation equilibria, we focus entirely on the separation equilibria in this section and leave analysis of the partial-pooling equilibria for the Appendix.

Proposition 4, below, describes the class of informative equilibria involving full separation of good and bad challengers. Unlike the other informative equilibria discussed in the no-agenda-setter game, the first three challenger separating equilibria involve disciplining the good incumbent into choosing the inappropriate policy. This is a very undesirable effect of political debate and shows that perverse outcomes can arise when the challenger stakes out a policy position before the incumbent. The fourth challenger separating equilibrium, however, involves only disciplining of the bad incumbent. Although this very desirable effect of political debate is similar to the positive effects described in the no-agenda-setter game, we will see that this equilibrium occurs for different reasons.

Proposition 4 Challenger Separation. Depending on the parameters, one of four generic informative equilibria may exist in which, w.l.o.g., the good challenger sends policy recommendation T and the bad challenger sends policy recommendation F .

1. **Fully Good-Incumbent Disciplining and Fully Bad-Incumbent Disciplining.** *There is an informative equilibrium in which the good and bad incum-*

bents implement T when $m = T$ and F when $m = F$, provided:

$$H(\beta_I, \beta_C) - H(0, 1) \geq b_G \quad (4.1-A)$$

$$H(\beta_I, \beta_C) - H(0, 1) \geq b_B \quad (4.1-B)$$

- 2. Partially Incumbent and Challenger Revealing, Fully Good-Incumbent Disciplining, and Bad-Incumbent Disciplining.** *There is an informative equilibrium in which: (i) the good incumbent implements T when $m = T$ and F when $m = F$, and (ii) the bad incumbent mixes over the policies with probability $q^* \in (0, 1)$ on T when $m = T$ and implements F when $m = F$, provided:*

$$H(\widehat{\beta}_I(q^*), \widehat{\beta}_C(q^*)) - H(0, 1) \geq b_G \quad (4.2-A)$$

$$H(\widehat{\beta}_I(0), \widehat{\beta}_C(0)) - H(0, 1) > b_B > H(\beta_I, \beta_C) - H(0, 1) \quad (4.2-B)$$

$$\beta_I \geq \frac{1}{2} \left(1 - \frac{q^*}{1 - q^*}\right) \quad (4.2-C)$$

where

$$\begin{aligned} \widehat{\beta}_I(q) &= \frac{\beta_I}{\beta_I + (1 - \beta_I)(\beta_C q + 1 - \beta_C)} \\ \widehat{\beta}_C(q) &= \frac{\beta_C[\beta_I + (1 - \beta_I)q]}{\beta_C[\beta_I + (1 - \beta_I)q] + 1 - \beta_C}. \end{aligned}$$

and q^* is implicitly given by

$$b_B = H(\widehat{\beta}_I(q^*), \beta_C(q^*)) - H(0, 1)$$

- 3. Partially Incumbent and Challenger Revealing and Fully Good-Incumbent Disciplining.** *There is an informative equilibrium in which: (i) the good in-*

cumbent implements T when $m = T$ and F when $m = F$, and (ii) the bad incumbent implements F after each policy suggestion, provided:

$$H(\widehat{\beta}_I, \widehat{\beta}_C) - H(0, 1) \geq b_G \quad (4.3-A)$$

$$b_B \geq H(\widehat{\beta}_I, \widehat{\beta}_C) - H(0, 1) \quad (4.3-B)$$

$$\beta_I \geq \frac{H(\widehat{\beta}_I, \widehat{\beta}_C) - H(0, 1)}{b_B + H(\widehat{\beta}_I, \widehat{\beta}_C) - H(0, 1)} \quad (4.3-C)$$

where

$$\widehat{\beta}_I = \frac{\beta_I}{\beta_I + (1 - \beta_I)(1 - \beta_C)} \quad \text{and} \quad \widehat{\beta}_C = \frac{\beta_C \beta_I}{\beta_C \beta_I + 1 - \beta_C}.$$

- 4. Bad-Incumbent Disciplining.** *There is an informative equilibrium in which:*
(i) the good incumbent implements T after each message, and (ii) the bad incumbent mixes over the policies with probability $q^* \in (0, 1)$ on T when $m = T$ and implements F when $m = F$, provided:

$$b_G \geq H(\widehat{\beta}_I^x(q^*), \widehat{\beta}_C^x(q^*)) - H(\widehat{\beta}_I^y(q^*), \widehat{\beta}_C^y(q^*)) \quad (4.4-A)$$

$$H(\widehat{\beta}_I^x(0), \widehat{\beta}_C^x(0)) - H(\widehat{\beta}_I^y(0), \widehat{\beta}_C^y(0)) > b_B > H(\widehat{\beta}_I^x(1), \widehat{\beta}_C^x(1)) - H(\widehat{\beta}_I^y(1), \widehat{\beta}_C^y(1)) \quad (4.4-B)$$

$$\frac{1 + q^*(b_G/b_B - 1)}{2 + q^*(b_G/b_B - 1)} > \beta_I > \frac{1}{2} \left(1 - \frac{q^*}{1 - q^*} \right) \quad (4.4-C)$$

where

$$\begin{aligned}
\widehat{\beta}_I^x(q) &= \frac{\beta_C \beta_I}{\beta_C \beta_I + \beta_C(1 - \beta_I)q + (1 - \beta_C)(1 - \beta_I)} \\
\widehat{\beta}_I^y(q) &= \frac{(1 - \beta_C)\beta_I}{(1 - \beta_C)\beta_I + \beta_C(1 - \beta_I)(1 - q)} \\
\widehat{\beta}_C^x(q) &= \frac{\beta_C[\beta_I + (1 - \beta_I)q]}{\beta_C[\beta_I + (1 - \beta_I)q] + (1 - \beta_C)(1 - \beta_I)} \\
\widehat{\beta}_C^y(q) &= \frac{\beta_C(1 - \beta_I)(1 - q)}{\beta_C(1 - \beta_I)(1 - q) + (1 - \beta_C)\beta_I}.
\end{aligned}$$

and q^* is implicitly given by

$$b_B = H(\widehat{\beta}_I^x(q^*), \beta_C^x(q^*)) - H(\widehat{\beta}_I^y(q^*), \beta_C^y(q^*))$$

In the first equilibrium above, Proposition 4.1, the challenger's policy recommendation never reveals the appropriate policy to the citizen (on the equilibrium path), but affects the policy choice of the good and bad incumbents vis-à-vis the babbling equilibrium. Here, the good and bad incumbents pool by always following the challenger's policy recommendation. Therefore, regardless of which policy the challenger recommends, the citizen always observes event x on the equilibrium path. His updated beliefs equal his initial beliefs and hence he votes for the incumbent with probability $H(\beta_I, \beta_C)$. Off the equilibrium path, the citizen can observe events x or y . If the challenger deviates, the citizen will observe event x and hence his updated beliefs are identical to those above. If the incumbent deviates, the citizen will observe event y . In this case, the citizen's voting rule cannot be tied down because the updated beliefs cannot be determined using Bayes's Rule. Therefore, let $\widehat{\beta}_I^x(y) = 0$ and $\widehat{\beta}_C^x(y) = 1$. That is, the citizen believes event y is generated by a bad incumbent when she deviates after a good challenger's policy recommendation.

The behavior of each politician in Proposition 4.1 is determined by the trade-

offs she faces between policy and winning election. Conditions (4.1-A) and (4.1-B) are the incentive compatibility constraints for the good and bad incumbents, respectively. If (4.1-A) is satisfied, the good incumbent's best response to any policy recommendation is to choose the policy that is recommended—she is willing to forgo her policy payoff for an increase in her expected reelection payoff by choosing the policy she does not prefer. By symmetry, the same is true for the bad incumbent if (4.1-B) is satisfied. Notice, it is a strict best response for an incumbent to follow a challenger's message when the recommendation is the incumbent's preferred policy. The good and bad challenger recommend their preferred policy because their policy recommendation determines which policy the incumbent implements, but does not affect on the citizen's voting rule.

As mentioned earlier, we find that political debate can negatively affect the citizen welfare when the challenger is the agenda-setter. Proposition 4.1. is a clear example. Here, the challenger gets control of the policy because incumbents always follow the challenger's recommendation. The good and bad incumbents do so because they care more for holding office than choosing their preferred policies (i.e., b_G and b_B sufficiently close to zero). Hence, political debate never informs the citizen of the appropriate policy (on the equilibrium path). Notice that the set of policy payoffs which support this equilibrium increases as the incumbent's reputation increases and as the challenger's reputation decreases. That is, the challenger can get control of the policy and implement her preferred policy when the incumbent is believed to be good while the challenger is believed to be bad. Under these circumstances, political debate disciplines the incumbent into choosing the inappropriate policy exactly when we would expect the incumbent to choose the appropriate policy were there no debate. Hence, debate harms the citizen.

In the second equilibrium, Proposition 4.2, the challenger's policy recommen-

dation partially informs the citizen of the appropriate policy and affects the policy choice of the good and bad incumbents vis-à-vis the babbling equilibrium. Here, the good incumbent always follows the challenger’s policy recommendation whereas the bad incumbent follows the policy recommendation if it is F , but otherwise strictly mixes over the two policies. The citizen observes events x and y on the equilibrium path because the good and bad challengers separate by recommending the appropriate and inappropriate policies respectively. If the citizen observes event x , he believes the incumbent and challenger are good with probability $\widehat{\beta}_I(q)$ and $\widehat{\beta}_C(q)$ respectively and votes for the incumbent with probability $H(\widehat{\beta}_I(q), \widehat{\beta}_C(q))$ where q is the probability the bad incumbent places on policy T when mixing. If he observes event y , he believes the incumbent and challenger are good with probability zero and one, respectively, and votes for the incumbent with probability $H(0, 1)$.

Just as in Proposition 4.1, the behavior of each politician in Proposition 4.2 is determined by the trade-offs she faces between policy and winning election. Conditions (4.2-A) and (4.2-B) are the incentive compatibility constraints for the good and bad incumbents, respectively. If (4.2-A) is satisfied, the good incumbent’s best response to any policy recommendation is to choose the policy that is recommended—she is willing to forgo her policy payoff for an increase in her expected reelection payoff by choosing the inappropriate policy. If (4.2-B) is satisfied, the bad incumbent’s best response to policy recommendation T is to strictly mix over the two policies—she is willing to forgo her policy payoff for an increase in her expected reelection payoff by choosing the appropriate policy, but as she increasingly plays the appropriate policy, her expected reelection payoff begins to fall so much so that she does not want to completely pool with the good incumbent. Again notice, it is a strict best response for an incumbent to follow a challenger’s message when the recommendation is the incumbent’s preferred policy. The good challenger has a strict incentive to send mes-

sage T because she strictly benefits from an increase in her expected election payoff and an increase in the frequency with which she receives her policy payoff. The bad challenger will send message F provided doing so satisfies her incentive compatibility constraint (4.2-C). That is, when she chooses message F , the increase in the frequency with which she receives her policy payoff must outweigh the decrease in her expected election payoff.

Proposition 4.2 is another example of political debate negatively affecting the citizen's welfare. Here, the good challenger recommends the appropriate policy and disciplines the bad incumbent into choosing the appropriate policy with some positive probability (strictly less than one) while the bad challenger recommends the inappropriate policy and disciplines the good incumbent into choosing the inappropriate policy with probability one. The former disciplining effect shows the possibility that this may increase the citizen's welfare. The latter disciplining effect shows the possibility that this may decrease the citizen's welfare. These effects exist when the good incumbent is more concerned about reputation than the bad incumbent—that is, when the good incumbent cares more for holding office relative to policy than does the bad incumbent (i.e., $b_B > b_G$). Moreover, these equilibria exist when the incumbent is likely to be good and the challenger is likely to be bad. Under these circumstances, it is the latter disciplining effect that wins. Just as in Proposition 4.1., political debate disciplines the incumbent into choosing the inappropriate policy exactly when we would expect the incumbent to choose the appropriate policy were there no debate. Hence, debate harms the citizen.

In the third equilibrium, Proposition 4.3, the challenger's policy recommendation partially informs the citizen of the appropriate policy, but only affects the policy choice of the good incumbent vis-à-vis the babbling equilibrium. Here, the good incumbent always follows the challenger's policy recommendation whereas the

bad incumbent follows the policy recommendation if it is F , but otherwise chooses the inappropriate policy (i.e., the bad incumbent always chooses the inappropriate policy regardless of the recommendation). This case can be thought of as the situation where the bad incumbent's mixing probability from Proposition 4.2 is equal to zero. The citizen observes events x and y on the equilibrium path because the good and bad challengers separate by recommending the appropriate and inappropriate policies respectively. If the citizen observes event x , he believes the incumbent and challenger are good with probability $\widehat{\beta}_I$ and $\widehat{\beta}_C$ respectively and votes for the incumbent with probability $H(\widehat{\beta}_I, \widehat{\beta}_C)$. If he observes event y , he believes the incumbent and challenger are good with probability zero and one, respectively, and votes for the incumbent with probability $H(0, 1)$.

Again, the behavior of each politician in Proposition 4.3 is determined by the trade-offs she faces between policy and winning election. Condition (4.3-A) is an incentive compatibility constraint for the good incumbent and plays the same role as (4.2-A) in Proposition 4.2. Condition (4.3-B) is an incentive compatibility constraint for the bad incumbent, but plays a different role from (4.2-B) in Proposition 4.2. Unlike (4.2-B), if (4.3-B) is satisfied, the bad incumbent's best response to policy recommendation T is to choose the inappropriate policy—she benefits more from the gain of her policy payoff than the loss she suffers from the decrease in her expected reelection payoff. Again notice, it is a strict best response for an incumbent to follow a challenger's message when the recommendation is the incumbent's preferred policy. The good challenger has a strict incentive to send message T because she strictly benefits from an increase in her expected election payoff and an increase in the frequency with which she receives her policy payoff. The bad challenger will send message F provided doing so satisfies her incentive compatibility constraint (4.3-C). That is, when she chooses message F , the increase in the frequency with which she

receives her policy payoff must outweigh the decrease in her expected election payoff.

Proposition 4.3 is a more extreme example of political debate’s negative effect captured in Proposition 4.2. Here, the good challenger’s recommendation has no disciplining effect on the bad incumbent. That is, political debate has no positive effect. It only disciplines the good incumbent into choosing the inappropriate policy. Again, debate harms the citizen.

In the fourth equilibrium, Proposition 4.4, the challenger’s policy recommendation partially informs the citizen of the appropriate policy, but only affects the policy choice of the bad incumbent vis-à-vis the babbling equilibrium. Here, the good incumbent follows the policy recommendation if it is T , but otherwise chooses the appropriate policy (i.e., the good incumbent always chooses the appropriate policy regardless of the recommendation), whereas the bad incumbent follows the policy recommendation if it is F , but otherwise strictly mixes over the two policies. The citizen observes events x and y on the equilibrium path because the good and bad challengers separate by recommending the appropriate and inappropriate policies respectively. If the citizen observes event x , he believes the incumbent and challenger are good with probability $\widehat{\beta}_I^x(q)$ and $\widehat{\beta}_C^x(q)$, respectively, and votes for the incumbent with probability $H(\widehat{\beta}_I^x(q), \widehat{\beta}_C^x(q))$ where q is the probability the bad incumbent places on policy T when mixing. If he observes event y , he believes the incumbent and challenger are good with probability $\widehat{\beta}_I^y(q)$ and $\widehat{\beta}_C^y(q)$, respectively, and votes for the incumbent with probability $H(\widehat{\beta}_I^y(q), \widehat{\beta}_C^y(q))$.

Again, the behavior of each politician in Proposition 4.4 is determined by the trade-offs she faces between policy and winning election. Conditions (4.4-A) and (4.4-B) are the incentive compatibility constraints for the good and bad incumbents, respectively. If (4.4-A) is satisfied, the good incumbent’s best response to any policy recommendation is to choose the policy she prefers—she benefits more from the

gain of her policy payoff than the loss she suffers from the decrease in her expected reelection payoff. If (4.4-B) is satisfied, the bad incumbent's best response to policy recommendation T is to strictly mix over the two policies—she is willing to forgo her policy payoff for an increase in her expected reelection payoff by choosing the appropriate policy, but as she increasingly plays the appropriate policy, her expected reelection payoff begins to fall so much so that she does not want to completely pool with the good incumbent. Again notice, it is a strict best response for an incumbent to follow a challenger's message when the recommendation is the incumbent's preferred policy. The good challenger will send message T provided doing so satisfies her incentive compatibility constraint, the left-hand side inequality in (4.4-C). The bad challenger will send message F provided doing so satisfies her incentive compatibility constraint, the right-hand side inequality in (4.4-C). That is, a challenger will recommend her preferred policy, provided the increase in the frequency with which she receives her policy payoff outweighs the decrease in her expected election payoff.

Unlike the challenger separating equilibria discussed thus far, Proposition 4.4 shows how political debate can improve the citizen's welfare when challengers recommend different policies. Here, the good challenger recommends the appropriate policy and disciplines the bad incumbent into choosing the appropriate policy with some positive probability. This effect exists when the bad incumbent is more concerned about reputation than the good incumbent—that is, when the bad incumbent cares more for holding office relative to policy than does the good incumbent (i.e., $b_G > b_B$). Moreover, these equilibria exist when the challenger is more likely to be good than the incumbent, i.e., when the citizen can trust the challenger more than the incumbent ($\beta_C > \beta_I$). Here political debate disciplines the incumbent into choosing the appropriate policy exactly when we would expect the incumbent to choose the inappropriate policy were there no debate. Hence, debate helps the citizen.

The reader may be curious as to why there is no equilibrium in which the opposite of Proposition 4.4 occurs, specifically, one in which the bad incumbent always follows the challenger's policy recommendation whereas the good incumbent follows the policy recommendation if it is T , but otherwise chooses the appropriate policy (i.e., the good incumbent always chooses the appropriate policy regardless of the recommendation). The citizen would clearly benefit from incumbents behaving in this manner. Unfortunately, no such equilibrium exists. To behave as proscribed above, the good incumbent must value implementing the appropriate policy over winning election. That is, she faces the following incentive compatibility constraint:

$$b_G \geq H(\widehat{\beta}_I^x, \widehat{\beta}_C^x) - H(\widehat{\beta}_I^y, \widehat{\beta}_C^y)$$

The bad incumbent must value winning election over implementing the inappropriate policy. That is, she faces the following incentive compatibility constraint:

$$H(\widehat{\beta}_I^x, \widehat{\beta}_C^x) - H(\widehat{\beta}_I^y, \widehat{\beta}_C^y) \geq b_B$$

Notice that because $b_B > 0$, we must have $H(\widehat{\beta}_I^x, \widehat{\beta}_C^x) - H(\widehat{\beta}_I^y, \widehat{\beta}_C^y) > 0$. Provided $H(\widehat{\beta}_I^x, \widehat{\beta}_C^x) - H(\widehat{\beta}_I^y, \widehat{\beta}_C^y) > 0$, the bad challenger has a strict best response to recommend the inappropriate policy as doing so increases the frequency she receives her policy payoff and increases her expected election payoff. But this implies that, regardless of the good incumbent's recommendation, $0 \geq H(\widehat{\beta}_I^x, \widehat{\beta}_C^x) - H(\widehat{\beta}_I^y, \widehat{\beta}_C^y)$. A contradiction. Given the behavior of the bad challenger and both incumbents, event y increases the incumbent's reputation and decreases the challenger's reputation while event x decreases the incumbent's reputation and increases the challenger's reputation.

Figures 2 through 5 depict all of the informative equilibria for the setting




Figure 2: Challenger as Agenda-Setter – Case 1: $\beta_I < 1/2$ and $\beta_I < \beta_C$

where the challenger is the agenda-setter under the assumption that the citizen's voting rule is the same linear function specified earlier. The regions labeled Proposition 5 correspond to the informative equilibria in which the challengers pool; those labeled Proposition 6 correspond to the informative equilibria in which the challengers partially-pool. As mentioned above, the interested reader can find discussion of these equilibria in the Appendix. Notice that some regions overlap, i.e., there are some regions with multiple informative equilibria.

1.4 Conclusion

We have presented a political agency model in which the challenger is a strategic player who engages the incumbent in a political debate prior to the election. When the citizen is unable to use the incumbent's policy choice as a means of screening good and

bad politicians, we have shown that political debate can provide the citizen with policy information which can be used to screen politicians. Under certain circumstances, the screening efforts of the citizen discipline bad incumbents into choosing appropriate policies and increase the citizen's welfare. However, under other circumstances, the screening efforts of the citizen discipline good incumbents into choosing inappropriate policies and decrease the citizen's welfare.

We find that the disciplining effects are sensitive to which politician, if any, is the agenda-setter. Moreover, we obtain that the welfare-increasing disciplinary effects arise when bad incumbents value winning election over implementing inappropriate policies and good politicians value implementing appropriate policies over winning election. Similarly, we obtain that the welfare-decreasing disciplinary effects arise when good incumbents value winning election over implementing appropriate policies and bad politicians value implementing inappropriate policies over winning election. Furthermore, we find that welfare-increasing disciplinary effects exist when the challenger has a better reputation than the incumbent and also when the incumbent is believed to be, on average, bad.

Appendix

No Agenda-Setter

Proof. Lemma 1.

Step 1. Suppose the good and bad incumbents pool on some mixed strategy. Furthermore, suppose the challenger sends message T with probability $\phi \in [0, 1]$.

Regardless of ϕ , the citizen may observe event x or event y . Let $h(x)$ be the probability the citizen reelects the incumbent when he observes event x . Let $h(y)$ be the probability he reelects the incumbent when he observes event y .

Consider the good incumbent's decision. If she chooses policy T , her expected payoff is $b_G + \phi h(x) + (1 - \phi)h(y)$. If she chooses policy F , her expected payoff is $\phi h(y) + (1 - \phi)h(x)$. As the good incumbent is mixing between T and F , it must be that she is indifferent between T and F :

$$b_G + (1 - 2\phi)[h(y) - h(x)] = 0 . \quad (\text{A1})$$

Consider the bad incumbent's decision. If she chooses policy T , her expected payoff is $\phi h(x) + (1 - \phi)h(y)$. If she choose policy F , her expected payoff is $b_B + \phi h(y) + (1 - \phi)h(x)$. As the bad incumbent is mixing between T and F , it must be that she is indifferent between T and F :

$$b_G - (1 - 2\phi)[h(y) - h(x)] = 0 . \quad (\text{A2})$$

Summing (A1) and (A2) yields $b_G + b_B = 0$ which is a contradiction as $b_G, b_B > 0$ by assumption.

Step 2. Suppose the incumbents pool on some pure strategy. Furthermore, suppose the challenger is strictly mixing over her messages.

In this situation, the citizen cannot update the incumbent's and challenger's initial reputations. As a consequence, the incumbent's policy choice does not affect her reelection probability. Therefore, each incumbent has strict incentive to choose her preferred policy. A contradiction.

Step 3. Suppose the incumbents pool on some pure strategy. Furthermore, suppose the challenger playing pure strategy T w.l.o.g.

Regardless of which pure strategy the incumbents pool on, the citizen will observe one event on the equilibrium path and the other event off the equilibrium path. On the equilibrium path, the citizen cannot update the politicians' initial

reputations and as a result reelects the incumbent with probability $H(\beta_I, \beta_C)$. Off the equilibrium, the citizen's beliefs regarding the politicians cannot be determined by Bayes's Rule and the citizen's reelection rule cannot be tied down. Therefore, let z be the probability the citizen reelects the incumbent off the equilibrium path.

The challenger will send message T if and only if her election probability is greater on the equilibrium path than off the equilibrium path, i.e., $z \geq H(\beta_I, \beta_C)$.

Now suppose the incumbents pool on T . In this case, the bad incumbent has a strict incentive to deviate because in doing so she gets her preferred policy and a higher reelection probability. A contradiction. Now suppose instead that the incumbents pool on F . In this case, the good incumbent has a strict incentive to deviate because in doing so she gets her preferred policy and a higher reelection probability. A contradiction. ■

Proof. Lemma 2.

Step 1. Suppose the good incumbent plays pure strategy F . From Lemma 1, we know the bad incumbent is either playing pure strategy T or is strictly mixing. Furthermore, suppose the challenger sends message T with probability $\phi \in [0, 1]$.

Regardless of ϕ , the citizen may observe event x or event y . Let $h(x)$ be the probability the citizen reelects the incumbent when he observes event x . Let $h(y)$ be the probability he reelects the incumbent when he observes event y . Notice, if $\phi \geq (\leq) \frac{1}{2}$ then $h(x) \leq (\geq) h(y)$.

Consider the bad incumbent's decision. If she chooses policy T , her expected payoff is $\phi h(x) + (1 - \phi)h(y)$. If she chooses policy F , her expected payoff is $b_B + \phi h(y) + (1 - \phi)h(x)$. As the bad incumbent is either playing pure strategy T or strictly mixing, it must be that she weakly prefers to choose policy T :

$$(1 - 2\phi)[h(y) - h(x)] \geq b_B .$$

This is a contradiction as the left-hand side is non-positive for all $\phi \in [0, 1]$.

Step 2. Suppose the good incumbent strictly mixes with probability $p \in (0, 1)$ on T and the bad incumbent strictly mixes with probability $q \in (0, 1)$ on F . Furthermore, suppose the challenger sends message T with probability $\phi \in [0, 1]$.

Regardless of ϕ , the citizen may observe event x or event y . Let $h(x)$ be the probability the citizen reelects the incumbent when he observes event x . Let $h(y)$ be the probability he reelects the incumbent when he observes event y .

As both the good and bad incumbents are strictly mixing, they must both be indifferent between choosing policies T and F . Hence, the good and bad incumbents' payoffs must respectively satisfy:

$$\begin{aligned} b_G + (1 - 2\phi)[h(y) - h(x)] &= 0 & \text{and} \\ b_G - (1 - 2\phi)[h(y) - h(x)] &= 0 . \end{aligned}$$

Summing these two conditions yields $b_G + b_B = 0$ which is a contradiction as $b_G, b_B > 0$ by assumption.

Step 3. Suppose the good incumbent strictly mixes with probability $p \in (0, 1)$ on T . From Step 2 above and Lemma 1, we know the bad incumbent is either playing pure strategy T or pure strategy F —w.l.o.g., suppose the bad incumbent plays T . Furthermore, suppose the challenger sends message T with probability $\phi \in [0, 1]$.

Regardless of ϕ , the citizen may observe event x or event y . Let $h(x)$ be the probability the citizen reelects the incumbent when he observes event x . Let $h(y)$ be the probability he reelects the incumbent when he observes event y . Notice, if $\phi \geq (\leq) \frac{1}{2}$ then $h(x) \geq (\leq) h(y)$.

Consider the good incumbent's decision. If she chooses policy T , her expected payoff is $b_G + \phi h(x) + (1 - \phi)h(y)$. If she chooses policy F , her expected payoff is

$\phi h(y) + (1 - \phi)h(x)$. As the good incumbent is strictly mixing, it must be that she is indifferent between T and F :

$$b_G + (1 - 2\phi)[h(y) - h(x)] = 0 .$$

This is a contradiction as the left-hand side is strictly positive for all $\phi \in [0, 1]$. ■

Challenger As Agenda-Setter

The following proposition describes the informative equilibria when challengers pool. As mentioned earlier, these equilibria are similar to those in Proposition 3.

Proposition 5 Challenger Pooling. Depending on the parameters, one of four generic informative equilibria may exist in which the good and bad challengers pool on policy suggestion T w.l.o.g.

1.a. Incumbent Revealing. *There is an informative equilibrium in which: (i) the good incumbent implements T after each policy suggestion, and (ii) the bad incumbent implements F after each policy suggestion, provided:*

$$b_G \geq H(1, \beta_C) - H(0, \beta_C) \tag{5.1.a-A}$$

$$b_B \geq H(1, \beta_C) - H(0, \beta_C) \tag{5.1.a-B}$$

$$\frac{1}{2} \geq \beta_I \tag{5.1.a-C}$$

1.b. Incumbent Revealing. *There is an informative equilibrium in which: (i) the good incumbent implements T when $m = T$ and F when $m = F$, and (ii) the*

bad incumbent implements F after each policy suggestion, provided:

$$H(1, \beta_C) - H(0, \beta_C) \geq b_G \quad (5.1.b-A)$$

$$b_B \geq H(1, \beta_C) - H(0, \beta_C) \quad (5.1.b-B)$$

$$\frac{H(1, \beta_C) - H(0, \beta_C)}{b_B + H(1, \beta_C) - H(0, \beta_C)} \geq \beta_I \quad (5.1.b-C)$$

2.a. Partially Incumbent Revealing and Bad-Incumbent Disciplining. *There is an informative equilibrium in which: (i) the good incumbent implements T after each policy suggestion, and (ii) the bad incumbent mixes over the policies with probability $q^* \in (0, 1)$ on T when $m = T$ and implements F when $m = F$, provided:*

$$b_G \geq H\left(\frac{\beta_I}{\beta_I + (1 - \beta_I)q^*}, \beta_C\right) - H(0, \beta_C) \quad (5.2.a-A)$$

$$H(1, \beta_C) - H(0, \beta_C) > b_B > H(\beta_I, \beta_C) - H(0, \beta_C) \quad (5.2.a-B)$$

$$\frac{1}{2} \left(1 - \frac{q^*}{1 - q^*}\right) \geq \beta_I \quad (5.2.a-C)$$

where q^* is implicitly given by

$$b_B = H(\widehat{\beta}_I(x|q^*), \beta_C) - H(0, \beta_C)$$

2.b. Partially Incumbent Revealing and Bad-Incumbent Disciplining. *There is an informative equilibrium in which: (i) the good incumbent implements T when $m = T$ and F when $m = F$, and (ii) the bad incumbent mixes over the policies with probability $q^* \in (0, 1)$ on T when $m = T$ and implements F when*

$m = F$, provided:

$$H\left(\frac{\beta_I}{\beta_I + (1 - \beta_I)q^*}, \beta_C\right) - H(0, \beta_C) \geq b_G \quad (5.2.b-A)$$

$$H(1, \beta_C) - H(0, \beta_C) > b_B > H(\beta_I, \beta_C) - H(0, \beta_C) \quad (5.2.b-B)$$

$$\frac{1}{2}\left(1 - \frac{q^*}{1 - q^*}\right) \geq \beta_I \quad (5.2.b-C)$$

where q^* is implicitly given by

$$b_B = H(\widehat{\beta}_I(x|q^*), \beta_C) - H(0, \beta_C)$$

The equilibria in Proposition 5 are analogous to the equilibria in Proposition 3 in that they are, broadly speaking, outcome equivalent. Specifically, the on-equilibrium-path behavior in Proposition 5.1.a and 5.1.b and the equilibrium behavior in Proposition 3.1 are the same—the good and bad challengers pool, w.l.o.g., on message T , the good incumbent chooses policy T , and the bad incumbent chooses policy F . Hence, the equilibria in Propositions 5.1.a and 5.1.b are incumbent revealing in the same manner as the equilibrium in Proposition 3.1. Similarly, the on-equilibrium-path behavior in Proposition 5.2.a and 5.2.b and the equilibrium behavior in Proposition 3.2 are the same—the good and bad challengers pool, w.l.o.g., on message T , the good incumbent chooses policy T , and the bad incumbent strictly mixes over the two policies. Hence, the equilibria in Propositions 5.2.a and 5.2.b are partially incumbent revealing and bad-incumbent disciplining in the same manner as the equilibrium in Proposition 3.2.

Although the equilibria in Proposition 5 are analogous to the equilibria in Proposition 3, they are noticeably different in that they require an additional condition—an incentive compatibility constraint for the good incumbent. The additional condi-

tion is necessary because, unlike the no-agenda-setter game where the incumbent has only one information set, in the challenger-as-agenda-setter game the incumbent has two information sets—she can observe recommendation T or F . Therefore, we need two conditions to tie down incumbent behavior instead of one.

When the challengers pool (w.l.o.g. on message T), only one of the incumbent’s two information sets will be on the equilibrium path. The B-suffix conditions are incentive compatibility constraints for the bad incumbent at the information set on the equilibrium path. Conditions (5.1.a-B) and (5.1.b-B) are identical to condition (3.1-B) in the no-agenda-setter game. Conditions (5.2.a-B) and (5.2.b-B) are identical to condition (3.2-B). As before, if (5.1.a-B) or (5.1.b-B) is satisfied, the bad incumbent’s best response to the challenger’s message is to separate from the good incumbent by choosing the inappropriate policy. If (5.2.a-B) or (5.2.b-B) is satisfied, the bad incumbent’s best response to the challenger’s message is to semi-separate from the good incumbent by strictly mixing over the two policies.

In the challenger-as-agenda-setter game, however, we also have to worry about incentives off the equilibrium path. The A-suffix conditions are incentive compatibility constraints for the good incumbent at the information set off the equilibrium path. If (5.1.a-A) or (5.2.a-A) is satisfied, the good incumbent’s best response to the challenger’s (off-equilibrium-path) message is to choose the appropriate policy—she benefits more from the gain of her policy payoff than the loss she suffers from the decrease in her expected reelection payoff. If (5.1.b-A) or (5.2.b-A) is satisfied, the good incumbent’s best response to the challenger’s (off-equilibrium-path) message is to choose the inappropriate policy—she is willing to forgo her policy payoff for an increase in her expected reelection payoff by choosing the inappropriate policy.

The good and bad challengers also face new incentives in the challenger-as-agenda-setter game. As before, the challenger’s policy recommendation has the po-

tential to influence the citizen's beliefs, but now also determines which of the incumbent's information sets is reached. The C-suffix conditions are incentive compatibility constraints for the challenger. Condition (5.1.a-C) is identical to condition (3.1-C) in the no-agenda-setter game. This is the result of the same behavior, on and off the equilibrium path, of the good and bad incumbents in Proposition 5.1.a. Condition (5.1.b-C) is similar to condition (3.1-C), but tighter. This is the result of the good incumbent pooling with the bad incumbent off the equilibrium path by choosing the inappropriate policy in Proposition 4.1.b. Deviating to message F is more (less) attractive for the bad (good) challenger because then the inappropriate policy is always chosen. Conditions (5.2.a-C) and (5.2.b-C) are identical to condition (3.2-C). This is the result of the bad incumbent's mixing behavior on the equilibrium path in Proposition 5.2.a and 5.2.b. The way the bad incumbent mixes counter acts the attractiveness or unattractiveness of deviating to message F .

Proposition 5 describes the affect political debate can have when the challenger is the agenda-setter provided the challengers pool on the policy recommendation. As before, political debate can inform the citizen of the appropriate policy if the incumbent is, on average, believed to be a bad politician. That is, debate can help citizens exactly when the citizen is more likely to receive the inappropriate policy from the incumbent were there no debate. Not surprisingly, as these equilibria are analogous to those in Proposition 3, we find roughly similar welfare effects from political debate. Generally speaking, everything is the same except for the situation described by Proposition 5.1.b. Just as in Proposition 3.1, bad politicians care more for inappropriate policies than they do for holding office tomorrow. Therefore political debate cannot discipline them into choosing the appropriate policy. However, good politicians care more for holding office tomorrow than they do for appropriate policies. This makes it more tempting for bad challengers to recommend inappropriate

policies. Bad challengers will refrain from recommending inappropriate policies provided the incumbent's reputation is sufficiently low. Therefore, the set of parameters under which the citizen can identify bad incumbents and improve his prospects for having a good politician in office next period in Proposition 5 is smaller than that in Proposition 3.

Figure 4 depicts the equilibria in Proposition 5 when we assume the citizen's voting function is the linear function specified earlier. Notice the region that supports the equilibria that are similar to Proposition 3.1 has shrunk. This reflects the tightening of the challenger's incentive compatibility constraint (5.1.b-C) as discussed above. Notice that equilibria involve challengers pooling do not overlap.

We now turn to the class of equilibria which involve the bad challenger partially pools with good challenger. These equilibria have no analog in the no-agenda-setter game and are described in the following proposition.

Proposition 6 Challenger Partial Pooling. Depending on the parameters, one generic informative equilibria may exist in which, w.l.o.g., the good challenger recommends policy T and the bad challenger strictly mixes over which policy to recommend with probability ϕ^* on T .

1. **Partially Incumbent Revealing and Fully Good-Incumbent Disciplining.** *There is an informative equilibrium in which: (i) the good incumbent implements T when $m = T$ and F when $m = F$, and (ii) the bad incumbents implement F after each policy suggestion, provided:*

$$H(\widehat{\beta}_I^x(\phi^*), \widehat{\beta}_C^x(\phi^*)) - H(\widehat{\beta}_I^y(\phi^*), \widehat{\beta}_C^y(\phi^*)) \geq b_G \quad (6.1-A)$$

$$b_B \geq H(\widehat{\beta}_I^x(\phi^*), \widehat{\beta}_C^x(\phi^*)) - H(\widehat{\beta}_I^y(\phi^*), \widehat{\beta}_C^y(\phi^*)) \quad (6.1-B)$$

$$\frac{H(\widehat{\beta}_I^x(0), \widehat{\beta}_C^x(0)) - H(0, 1)}{b_B + H(\widehat{\beta}_I^x(0), \widehat{\beta}_C^x(0)) - H(0, 1)} \geq \beta_I \geq \frac{H(1, \beta_C) - H(0, \beta_C)}{b_B + H(1, \beta_C) - H(0, \beta_C)} \quad (6.1-C)$$

where

$$\begin{aligned}\widehat{\beta}_I^x(\phi) &= \frac{\beta_I}{\beta_I + (1 - \beta_I)(1 - \beta_C)(1 - \phi)} \\ \widehat{\beta}_I^y(\phi) &= 0 \\ \widehat{\beta}_C^x(\phi) &= \frac{\beta_C\beta_I}{\beta_C\beta_I + (1 - \beta_C)\beta_I + (1 - \beta_C)(1 - \beta_I)(1 - \phi)} \\ \widehat{\beta}_C^y(\phi) &= \frac{\beta_C(1 - \beta_I)}{\beta_C(1 - \beta_I) + (1 - \beta_C)(1 - \beta_I)(1 - \phi)}\end{aligned}$$

and ϕ^* is implicitly given by

$$\beta_I = \frac{H(\widehat{\beta}_I^x(\phi^*), \widehat{\beta}_C^x(\phi^*)) - H(\widehat{\beta}_I^y(\phi^*), \widehat{\beta}_C^y(\phi^*))}{b_B + H(\widehat{\beta}_I^x(\phi^*), \widehat{\beta}_C^x(\phi^*)) - H(\widehat{\beta}_I^y(\phi^*), \widehat{\beta}_C^y(\phi^*))}$$

In the equilibrium above, Proposition 6.1, the challenger's recommendation partially informs the citizen of the appropriate policy and affects the policy choice of the good incumbent vis-à-vis the babbling equilibrium. Just as in Proposition 4.3, the good incumbent always follows the challenger's policy recommendation whereas the bad incumbent follows the policy recommendation if it is F , but otherwise chooses the inappropriate policy (i.e., the bad incumbent always chooses the inappropriate policy regardless of the recommendation). The key difference is that the bad challenger semi-separates from the good challenger. Therefore, the citizen observes events x and y on the equilibrium path as before. If the citizen observes event x , he believes the incumbent and challenger are good with probability $\widehat{\beta}_I^x(\phi)$ and $\widehat{\beta}_C^x(\phi)$ respectively and votes for the incumbent with probability $H(\widehat{\beta}_I^x(\phi), \widehat{\beta}_C^x(\phi))$ where ϕ is the probability the bad challenger places on message T when mixing. If the citizen observes event y , he believes the incumbent and challenger are good with probability one and $\widehat{\beta}_C^y(\phi)$ respectively and votes for the incumbent with probability $H(0, \widehat{\beta}_C^y(\phi))$.

Again, the behavior of each politician in Proposition 6.1 is determined by the trade-offs she faces between policy and winning election. Condition (6.1-A) is an incentive compatibility constraint for the good incumbent and plays the same role as (4.3-A) in Proposition 4.3. Condition (6.1-B) is an incentive compatibility constraint for the bad incumbent and plays the same role as (4.3-B). Just as before, notice that it is a strict best response for an incumbent to follow a challenger's message when the recommendation is the incumbent's preferred policy. The good challenger has a strict incentive to send message T because she strictly benefits from an increase in her expected election payoff and an increase in the frequency with which she receives her policy payoff. The bad challenger will mix over the two messages provided doing so satisfies her incentive compatibility constraint (6.1-C). That is, when she chooses message F , the decrease in her expected election payoff outweighs the increase in the frequency with which she receives her policy payoff. However, when she chooses message T , the decrease in the frequency with which she receives her policy payoff outweighs the increase in the expected election payoff. Hence, she strictly mixes.

Given the similarity of this equilibrium with that in Proposition 4.3, it is not surprising that political debate negatively affects the citizen's welfare in this case. Here, when the bad challenger recommends the inappropriate policy, the good incumbent is disciplined into choosing the inappropriate policy. As the bad challenger is strictly mixing, this undesirable disciplining effect occurs with less frequency than in Proposition 4.3. However, the result is the same. Political debate disciplines the incumbent into choosing the inappropriate policy exactly when we would expect the incumbent to choose the appropriate policy were there no debate. Hence, debate harms the citizen. The circumstances under which this equilibrium exists is also similar to Proposition 4.3. Specifically, this effect exists when the good incumbent is more concerned about reputation than the bad incumbent—that is, when the good

incumbent cares more for holding office relative to policy than does the bad incumbent (i.e., $b_B > b_G$). Moreover, these equilibria exist when the incumbent is more likely to be good than the challenger, i.e., when the challenger is trustworthy than the incumbent.

References

- Aragones, Enriqueta, Thomas Palfrey and Andrew Postlewaite. "Campaign Rhetoric: A Model of Reputation." mimeo (2004).
- Austen-Smith, David. "Credible Debate Equilibria." *Social Choice and Welfare* 7 (1990): 75-93.
- Banks, Jeffrey and Rangarajan Sundaram. "Optimal Retention in Agency Problems." *Journal of Economic Theory* 82 (1998): 293-323.
- Barro, Robert. "The Control of Politicians: An Economic Model." *Public Choice* 14 (1973): 19-42.
- Bohman, James. *Public Deliberation: Pluralism, Complexity, and Democracy*. Cambridge, MA: MIT Press, 1996.
- Coate, Stephen and Stephen Morris. "On the Form of Transfers to Special Interests." *Journal of Political Economy* 103 (1995): 1210-1235.
- Elster, Jon. "Deliberation and Constitution Making," in Jon Elster (ed.) *Deliberative Democracy*. Cambridge University Press: New York, 1998.
- Fearon, James. "Deliberation as Discussion," in Jon Elster (ed.) *Deliberative Democracy*. Cambridge University Press: New York, 1998.
- Ferejohn, John. "Incumbent Performance and Electoral Control." *Public Choice* 50 (1986): 5-25.
- Gutmann, Amy and Dennis Thompson. *Democracy and Disagreement: Why Moral Conflict Cannot Be Avoided In Politics and What Can Be Done About It*. Cambridge, MA: Harvard University Press, 1996.

Habermas, Jürgen. *Between Facts and Norms: Contributions to a Discourse Theory of Law and Democracy*. Cambridge, MA: MIT Press, 1996.

Mackie, Gerry. “All Men Are Liars: Is Democracy Meaningless?,” in Jon Elster (ed.) *Deliberative Democracy*. Cambridge University Press: New York, 1998.

Maskin, Eric and Jean Tirole. “The Politician and the Judge: Accountability in Government.” *American Economic Review* forthcoming.

Morris, Stephen. “Political Correctness.” *Journal of Political Economy* 109 (2001): 231-265.

Stokes, Susan. “Pathologies of Deliberation,” in Jon Elster (ed.) *Deliberative Democracy*. Cambridge University Press: New York, 1998.

Stasavage, David. “Public versus Private Deliberation in a Representative Democracy.” mimeo (2004).

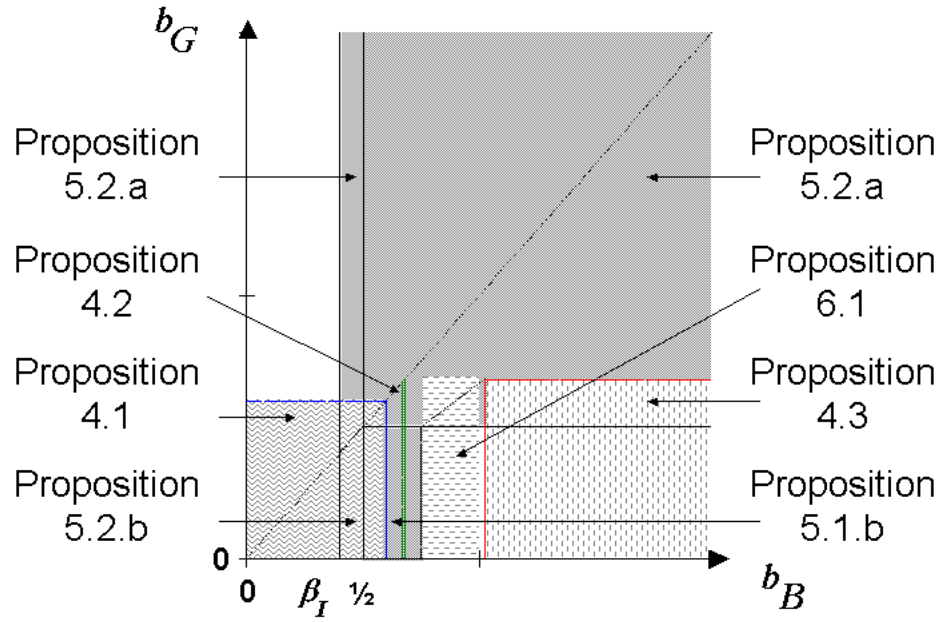


Figure 3: Challenger as Agenda-Setter – Case 2: $\beta_I < 1/2$ and $\beta_C < \beta_I$




Figure 4: Challenger as Agenda-Setter – Case 3: $1/2 < \beta_I$ and $\beta_I < \beta_C$




Figure 5: Challenger as Agenda-Setter – Case 4: $1/2 < \beta_I$ and $\beta_C < \beta_I$




Figure 6: Challenger as Agenda-Setter ($\beta < 1/2$)

Chapter 2

Flip-floppers versus Hard-liners

Keywords: flip-flopping, political agency, opportunism, elections, learning, signaling

JEL classification: D72, D82, D83

Abstract

Why do some candidates for office initially support one position on an election-year issue, only to later switch support to a contradictory position? Candidates who flip-flop in this manner would have us believe that new information has emerged which justifies a change in policy. Their political opponents say that flip-flopping has nothing to do with the merits of policies. Instead, they claim that candidates who flip-flop do so because it improves their prospects of winning the election. Their view suggests that flip-flopping signals political opportunism. This intuition is not captured by any existing political-agency model of flip-flopping. Indeed, a paper by Mukand and Majumdar shows the opposite: flip-flopping can signal that a candidate is of low ability. Assuming that citizens want to elect high-ability candidates, opportunistic politicians would never flip-flop. In this paper, we present a political-agency model which unifies these two divergent views of flip-flopping.

2.1 Introduction

Citizens are typically suspicious of a politician's policy motives: does a candidate advocate the policy she sincerely believes is best for society or the policy she believes is best for winning the election? Nowhere is this mistrust more evident than when a politician initially announces support for one policy, but later changes her support to another policy in the face of contradictory public opinion.

When a candidate flip-flops on an issue in this manner, she may be doing so sincerely or opportunistically. Suppose new information arrives which suggests the candidate's initial position is not in the public interest and that the other policy is more likely to be beneficial. One possibility is that the candidate updates on this information and learns that the other policy is best for society—she flip-flops because she honestly believes in the other position. Another possibility is that the candidate updates on this information, but continues to believe that her initial position is best for society—she flip-flops, however, because she believes the other position will improve her election prospects. The latter possibility leads to the conventional wisdom view on policy flip-flops: policy flip-flops signal that a politician is an opportunist, i.e., the politician is more concerned with getting elected than the appropriateness of the policies she espouses.

When a candidate maintains her initial position in the face of contradictory public opinion, however, she may also be doing so sincerely or opportunistically. Again, suppose that new information arrives which suggests the candidate's initial position is not in the public interest and that the other policy is more likely to be beneficial. One possibility is that the candidate updates on this information, but continues to believe that her initial position is best for society—she maintains her initial position because she honestly believes in it. Another possibility is that the candidate updates on this information and learns that the other policy is best

for society—she maintains her initial position, however, because she believes that if she flip-flops she will look incompetent and hurt her election prospects. The latter possibility leads to the Majumdar-Mukand (2004) view on policy flip-flops: policy flip-flops signal that a politician is incompetent, i.e., the politician is of low ability in identifying policies which benefit society. Provided citizens want to elect high-ability candidate, opportunists would never flip-flop.

The conventional wisdom understanding of flip-flopping and Majumdar-Mukand understanding of flip-flopping present two divergent views on opportunistic political behavior. This paper unifies these two views. We show that opportunistic flip-flopping can occur when citizens care greatly for immediate policy issues, while opportunistic policy maintenance can occur when citizens care greatly for future policy issues. The first section of this paper constructs the model. The second section analyzes how politicians choose policy positions. The third section concludes.

2.2 The Model

We construct a model in which an incumbent politician and challenger compete for elected office. The winning candidate must decide two policy issues: an election-year issue to which she must commit and an unforeseen future issue to which she cannot commit. The incumbent and challenger differ in their ability to identify the policy that best serves the public interest and in their willingness to sacrifice the public good for their desire to hold office. Each candidate observes her own private signal regarding which of two election-year policy positions benefits the citizen and then announces her initial position. The candidates and the citizen then observe a public signal regarding which election-year policy position benefits the citizen. Candidates update on this information and commit to a revised election-year policy position:

they can maintain their initial position or flip-flop. After candidates have revised their positions, the citizen elects the politician he believes is best based not only on the election-year issue, but also on the prospect of coping well with the future issue. The winning candidate implements her revised policy position and decides a new policy decision according to a new private signal.

Election-Year Policies. There are two election-year policies to which candidates can commit: l and r . There are also two states of the world: λ and ρ . Policy l benefits the citizen if and only if the state is λ while policy r benefits the citizen if and only if the state is ρ . At the start of the game, Nature chooses the state and determines which policy benefits the citizen. The incumbent and challenger accurately observe Nature's choice with probability q_I and q_C respectively. In other words, each candidate receives her own noisy private signal of the appropriate policy—we denote the incumbent's signal s_I and challenger's signal s_C . The citizen, however, knows only the prior that each state is equally likely and hold beliefs regarding the candidates' ability to accurately identify appropriate policies.

After observing their private signals, the candidates simultaneously adopt an initial policy position on the election-year issue. That is, the incumbent and challenger costlessly announce support for either policy l or r . After staking out their initial positions, additional information regarding which policy is appropriate becomes available in the form of a noisy public signal s_P which is correct with probability π . Candidates update on this new information and commit to a revised policy position on the election-year issue. They can maintain their initial position or switch support to the other position. The citizen updates on the public signal and the behavior of the politicians and then votes for the candidate who he believes is best based on two criteria: (i) the likelihood that the candidate's election-year policy position is beneficial, and (ii) the likelihood that the candidate will choose the appropriate policy on

the future issue.

Future Policies. After implementing her revised election-year policy position, the candidate that wins the election must decide a new issue that was unforeseen at the time of balloting. She alone receives a private signal regarding which of two policies is beneficial for the citizen and then implements a policy.

Politicians and the Citizen. All politicians prefer to implement appropriate policies, *ceteris paribus*. They can differ, however, in two dimensions: (i) their ability to identify appropriate policies, and (ii) their willingness to forgo the appropriate policy to improve their prospects of winning the election.

With respect to their ability to identify appropriate policies, politicians are one of two possible types: *high ability* or *low ability*. High-ability candidates, regardless of being an incumbent or challenger, receive private signals which correctly reveal the appropriate policy with probability q_H . Similarly, low-ability candidates receive private signals which correctly reveal the appropriate policy with probability q_L . We assume $1 > q_H > q_L > \frac{1}{2}$. In other words, both high-ability and low-ability politicians are better informed than the citizen regarding the appropriate policy, with the high-ability politician being better informed than the low-ability politician. No politician, however, is fully informed. It is common knowledge that β is the proportion of high-ability politicians, i.e., a politician is of high ability with probability β .

With respect to their willingness to forgo the appropriate policy to improve their prospects of winning the election, politicians are also one of two possible types: *sincere-myopic* or *opportunistic*. A sincere-myopic candidate does not place any value on winning the election and instead always implements the policy she believes is most likely to benefit the citizen. An opportunistic candidate only values holding office and therefore always acts to maximize her approval rating, i.e., what the citizen thinks

of her. It is common knowledge that θ is the proportion of opportunistic politicians, i.e., a politician is opportunistic with probability θ .

As mentioned earlier, the citizen has two desires: (i) an appropriate policy on the election-year issue, and (ii) an appropriate policy on the future issue. If the election-year issue is resolved appropriately then the citizen receives a normalized payoff of α . If the future issue is resolved appropriately then the citizen receives a normalized payoff of $1 - \alpha$. Issues that are resolved inappropriately provide the citizen with zero benefit.

The citizen observes the public signal regarding which election-year policy position is beneficial and each candidate's behavior, i.e., each candidate's initial election-year position and whether she maintained her position or flip-flopped. Using this information, the citizen updates his beliefs regarding each candidate's revised position on the election-year issue and her ability to address the future issue well. The citizen then forms the following approval rating for each candidate based on his expected payoff from electing that candidate:

$$\alpha \hat{p}_j + (1 - \alpha) \hat{q}_j \tag{1}$$

where \hat{p}_j is the citizen's update belief that candidate j 's revised policy position is appropriate and \hat{q}_j is the citizen's updated expectation that candidate j will handle the future issue appropriately which is simply his updated expectation of candidate j 's ability to identify appropriate policies. Although α and $1 - \alpha$, are respectively, the citizen's normalized payoffs from having the election-year issue and future issue resolved appropriately, they can also be interpreted as the weights the citizen respectively places on the two issues.

The Game and Definition of Equilibrium. This model defines a game between

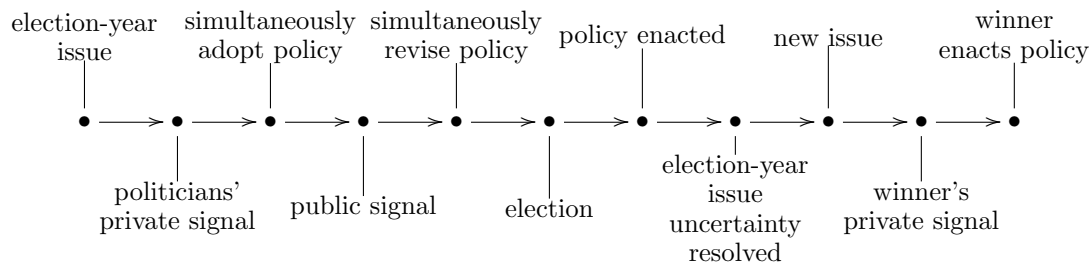


Figure 1: Time Line

the incumbent, challenger, and citizen. At the beginning of the game, Nature chooses the incumbent's type and the challenger's type. Nature then determines which policy is appropriate for the election-year issue. As mentioned above, each candidate receives her own noisy private signal regarding which policy is appropriate. The citizen knows only the prior that each policy is equally likely to be appropriate. Furthermore, the citizen is uncertain about each candidate's type. The incumbent and challenger are similarly uninformed about their opponent's type. That is, each candidate holds the same belief as the citizen regarding her opponent's type. After Nature has moved, the incumbent and challenger simultaneously announce their initial policy positions on the election-year issue. A public signal regarding which policy is appropriate follows. The candidates update on this information and commit to a revised position. The citizen observes the public signal and each candidate's behavior, and updates his beliefs. The election follows. The winner enacts her revised policy position and the uncertainty regarding election-year issue is resolved. A new issue emerges. The winner gets new private signal regarding which policy is best for this issue and enacts the policy she believes is most likely to benefit the citizen. Figure 1 presents the game's time line.

In this model, a Perfect Bayesian Nash Equilibrium (PBNE) consists of: (1) a strategy for each candidate, (2) each candidate's beliefs, (3) a strategy for the citizen, and (4) the citizen's beliefs. First, each candidate's strategy must be optimal

given her beliefs, her opponent's strategy, and the beliefs and strategy of the citizen. Second, each candidate's beliefs must be consistent with her opponent's strategy, i.e., beliefs must be formed according to Bayes's Rule where possible. Third, the citizen's strategy must be optimal given his beliefs. Fourth, the citizen's beliefs must be consistent with each candidate's strategy.

2.3 Policy Positioning

Throughout this paper, we assume that each candidate truthfully follows her private signal when announcing her initial policy position on the election-year issue, regardless of her type.¹ We focus on the interesting case in which the candidates' private signals agree with each other, but disagree with the public signal. Here, the scope for flip-flopping is greater as both candidates have the opportunity to do so. Without loss of generality, suppose that each candidate receives private signal r and that the subsequent public signal is l , i.e., $(s_I, s_C, s_P) = (r, r, l)$.

When a candidate is faced with the decision to revise her initial policy, she knows the realization of the private and public signals and the probability her private signal and the public signal are correct. She does not know, however, the probability that her opponent's signal is correct because she does not know her opponent's ability to identify appropriate policies. Instead, she had the prior belief that her opponent is of high ability with probability β and of low ability with probability $1 - \beta$, and ex-ante expects her opponent's signal is correct with probability $\beta q_H + (1 - \beta)q_L$.

As mentioned earlier, a sincere-myopic politician always chooses the policy she believes is best for the citizen. She updates on the above information and commits to the policy that she now believes is most likely to be appropriate. In other words,

¹In a more general model, we would check the validity of this assumption. Checking this assumption, however, is outside the scope of this paper.

when a sincere-myopic politician revises her initial policy position, she follows her posterior and therefore behaves efficiently with respect to her information. Given the triple $(s_I, s_C, s_P) = (r, r, l)$, a k -ability candidate knows her signal of r is correct with probability q_k , her opponent's signal of r is correct with probability $\beta q_H + (1 - \beta)q_L$, and the public signal of l is correct with probability π . Hence, she forms the following posterior that the state is ρ :

$$\Pr[\rho|r_k, r_A, l_\pi] = \frac{q_k(\beta q_H + (1 - \beta)q_L)(1 - \pi)}{q_k(\beta q_H + (1 - \beta)q_L)(1 - \pi) + (1 - q_k)(\beta(1 - q_H) + (1 - \beta)(1 - q_L))\pi} \quad (2)$$

where r_k is a signal r of q_k quality, r_A is a signal r of average quality (i.e., $\beta q_H + (1 - \beta)q_L$), and l_π is a public signal l of π quality. If this expression is equal to or greater than $\frac{1}{2}$, a k -ability sincere-myopic candidate will maintain her initial position r ; if this expression is less than $\frac{1}{2}$, a k -ability sincere-myopic candidate will switch her support to l and commit a flip-flop.

Figure 2 graphs expression (2) for both the high-ability and low-ability politicians as the informativeness of the public signal, π , increases. When π is close to $\frac{1}{2}$, the contradictory public signal is not very informative. Therefore, both the high-ability and low-ability politicians continue to believe in their own private signals—i.e., their posteriors on ρ are greater than $\frac{1}{2}$. As π increases, the informativeness of the contradictory public signal increases. Therefore, both the high-ability and low-ability politicians begin to question their own private signals—i.e., their posteriors on ρ begin to fall. Notice that the low-ability politician's posterior on ρ falls below $\frac{1}{2}$ before the high-ability politician's posterior on ρ . The intuition is that a low-ability politician is less confident in her own signal than a high-ability politician is in her own signal and, therefore, a low-ability politician has a lower posterior that her initial position is correct, *ceteris paribus*.

Throughout this paper, we restrict attention to parameter values such that,

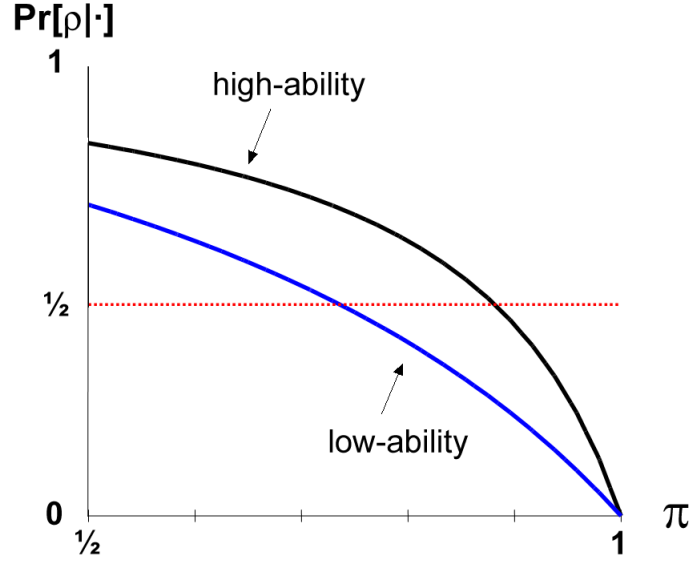


Figure 2: High-ability and Low-ability Politicians' Posteriors on ρ

when the incumbent and challenger take the same initial positions and are subsequently faced with a contradictory public signal, a high-ability politician continues to believe that her private signal is correct, but a low-ability politician now believes that her private signal is wrong. Graphically, we are restricting attention to values of π in Figure 1 where the high-ability politician's posterior is greater than $\frac{1}{2}$ the low-ability politician's posterior is less than $\frac{1}{2}$. We can represent this assumption by the following two inequalities:

$$\frac{q_H}{1 - q_H} \times \frac{\beta q_H + (1 - \beta)q_L}{\beta(1 - q_H) + (1 - \beta)(1 - q_L)} \geq \frac{\pi}{1 - \pi} \quad (\text{E1})$$

$$\frac{\pi}{1 - \pi} \geq \frac{q_L}{1 - q_L} \times \frac{\beta q_H + (1 - \beta)q_L}{\beta(1 - q_H) + (1 - \beta)(1 - q_L)} \quad (\text{E2})$$

Condition (E1) states that the contradictory public signal is not strong enough to undermine the high-ability politician's confidence in her initial position—the product of the odds ratio that her private signal is correct and the odds ratio that her expected opponent's private signal is correct is greater than the odds ratio that the public signal

is correct. Condition (E2) states that the public signal is strong enough to undermine the low-ability politician’s confidence in her initial position—the odds ratio that the public signal is correct is greater than the product of the odds ratio that her private signal is correct and the odds ratio that her expected opponent’s private signal is correct. Therefore, given (E1) and (E2), a high-ability sincere-myopic candidate maintains her initial position and a low-ability sincere-myopic candidate flip-flops.

Unlike sincere-myopic politicians, opportunists do not base their decision to maintain their initial position or flip-flop on the appropriateness of each policy. Instead, they act to maximize their expected approval rating. Given (E1) and (E2), we know that a high-ability politician’s posterior on ρ recommends maintaining her initial position and that a low-ability politician’s posterior on ρ recommends flip-flopping. Hence, opportunists can behave inefficiently in two ways: (i) a high-ability opportunist may inefficiently flip-flop to a policy that she believes is not in the citizen’s interest because it maximizes her expected approval rating, and (ii) a low-ability opportunist may inefficiently maintain a policy that she believes is not in the citizen’s interest because it maximizes her expected approval rating.

Lemma 1 establishes that, at most, only one type of opportunist can behave inefficiently with full probability in equilibrium.

Lemma 1 (Opportunist Behavior)

1. *If high-ability opportunists inefficiently flip-flop, then low-ability opportunists efficiently flip-flop.*
2. *If low-ability opportunists inefficiently maintain their initial position, then high-ability opportunists efficiently maintain their initial position.*

To prove the first part of Lemma 1, suppose both types of opportunists flip-flop. We need to show that a low-ability opportunist’s expected approval rating from flip-

flopping is greater than that from maintaining her initial policy whenever the high-ability opportunist’s expected approval rating from flip-flopping is greater than that from maintaining her initial position. A candidate’s expected approval rating is the convex combination of her expectation of the citizen’s posterior belief regarding her ability to address the future issue appropriately and her expectation of the citizen’s posterior that her revised position is appropriate. To understand the incentives that opportunists face, we examine each component of the candidate’s expected approval rating.

A candidate’s expectation of the citizen’s posterior belief regarding her ability is simply the the citizen’s posterior belief regarding her ability. When the citizen observes a candidate flip-flopping, he knows that the candidate is either a low-ability politician or a high-ability opportunist. His updated belief regarding the “average” ability of a candidate who has flip-flopped is:

$$q_{\underline{A}} = \frac{1 - \beta}{1 - \beta(1 - \theta)} q_L + \frac{\beta\theta}{1 - \beta(1 - \theta)} q_H \quad .$$

He believes that this candidate’s private signal r is of “average” quality which we denote $r_{\underline{A}}$. When the citizen observes a candidate maintaining her initial position, he knows that the candidate is a high-ability sincere-myopic politician. His updated belief regarding the ability of a candidate who has flip-flopped is q_H . He knows that this candidate’s private signal r is of high quality which we denote r_H .

A candidate’s expectation of the citizen’s posterior that her revised position is appropriate is more complicated because it depends on whether the candidate’s opponent flip-flops or maintains her initial position. If a candidate flip-flops to l , her opponent may also flip-flop in which case the citizen’s posterior on λ is $\Pr[\lambda|r_{\underline{A}}, r_{\underline{A}}, \lambda_{\pi}]$ or her opponent may maintain her initial position r in which case the citizen’s pos-

terior on λ is $\Pr[\lambda|r_{\underline{A}}, r_H, \lambda_\pi]$. If a candidate maintains r , her opponent may flip-flop in which case the citizen's posterior on ρ is $\Pr[\rho|r_H, r_{\underline{A}}, \lambda_\pi]$ or her opponent may also maintain r in which case the citizen's posterior on ρ is $\Pr[\lambda|r_H, r_H, \lambda_\pi]$. Although all politicians know that high-ability sincere-myopic politicians maintain their initial positions while high-ability opportunists and low-ability politicians flip-flop, high-ability and low-ability politicians will differ in their beliefs regarding their opponent's ability because private signals are correlated. We denote a k -ability candidate's posterior that her opponent is of high ability as $\hat{\beta}_k$. Hence, a k -ability candidate's expectation of the citizen's posterior on λ when she flip-flops is $(1 - \hat{\beta}_k(1 - \theta))\Pr[\lambda|r_{\underline{A}}, r_{\underline{A}}, l_\pi] + \hat{\beta}_k(1 - \theta)\Pr[\lambda|r_{\underline{A}}, r_H, l_\pi]$; her expectation of the citizen's posterior on ρ when she maintains her initial position is $(1 - \hat{\beta}_k(1 - \theta))\Pr[\rho|r_H, r_{\underline{A}}, l_\pi] + \hat{\beta}_k(1 - \theta)\Pr[\rho|r_H, r_H, l_\pi]$.

Given these components of a candidate's expected approval rating, the incentive compatibility constraint which guarantees that a k -ability opportunist will flip-flop when it is common knowledge that all opportunists flip-flop is:

$$\begin{aligned} & \alpha \left[(1 - \hat{\beta}_k(1 - \theta))\Pr[\lambda|r_{\underline{A}}, r_{\underline{A}}, l_\pi] + \hat{\beta}_k(1 - \theta)\Pr[\lambda|r_{\underline{A}}, r_H, l_\pi] \right] + (1 - \alpha)q_{\underline{A}} \\ & > \alpha \left[(1 - \hat{\beta}_k(1 - \theta))\Pr[\rho|r_H, r_{\underline{A}}, l_\pi] + \hat{\beta}_k(1 - \theta)\Pr[\rho|r_H, r_H, l_\pi] \right] + (1 - \alpha)q_H \quad (\text{K}) \end{aligned}$$

The LHS of (K) is a k -ability politician's expected approval rating from switching support to policy l ; the RHS of (K) is a k -ability politician's expected approval rating from maintaining support for policy r . The proof of the first part follows trivially from two observations: (i) the LHS of (K) decreases in $\hat{\beta}_k$ while the RHS of (K) increases in $\hat{\beta}_k$, and (ii) $\hat{\beta}_H > \hat{\beta}_L$. Hence, if flip-flopping to policy l is optimal for the high-ability opportunist, then flip-flopping must also be optimal for a low-ability opportunist.

To prove the second part of Lemma 1, suppose both types of opportunists maintain their initial positions. We need to show that a high-ability opportunist's

expected approval rating from maintaining her initial policy is greater than that from flip-flopping whenever the low-ability opportunist’s expected approval rating from maintaining her initial position is greater than that from flip-flopping. Recall that a candidate’s expected approval rating is the convex combination of her expectation of the citizen’s posterior belief regarding her ability to address the future issue appropriately and her expectation of the citizen’s posterior that her revised position is appropriate. We again examine each component of the candidate’s expected approval rating.

As mentioned above, a candidate’s expectation of the citizen’s posterior belief regarding her ability is simply the the citizen’s posterior belief regarding her ability. When the citizen observes a candidate maintaining her initial position, he knows that the candidate is either a high-ability politician or a low-ability opportunist. His updated belief regarding the “average” ability of a candidate who has maintained her position is:

$$q_A = \frac{(1 - \beta)\theta}{(1 - \beta)\theta + \beta} q_L + \frac{\beta}{(1 - \beta)\theta + \beta} q_H \quad .$$

He believes that this candidate’s private signal r is of “average” quality which we denote $r_{\bar{A}}$. When the citizen observes a candidate flip-flopping, he knows that the candidate is a low-ability sincere-myopic politician. His updated belief regarding the ability of a candidate who has flip-flopped is q_L . He knows that this candidate’s private signal r is of low quality which we denote r_L .

Again, a candidate’s expectation of the citizen’s posterior that her revised position is appropriate depends on whether the candidate’s opponent maintains her initial position or flip-flops. If a candidate maintains r , her opponent may also maintain r in which case the citizen’s posterior on ρ is $\Pr[\rho|r_{\bar{A}}, r_{\bar{A}}, l_\pi]$ or her opponent may flip-flop in which case the citizen’s posterior on ρ is $\Pr[\rho|r_{\bar{A}}, r_H, l_\pi]$. If a candidate flip-flops to l , her opponent maintain her initial position r in which case the citizen’s posterior

on λ is $\Pr[\lambda|r_L, r_{\bar{A}}, l_\pi]$ or her opponent may also flip-flop in which case the citizen's posterior on λ is $\Pr[\lambda|r_L, r_L, l_\pi]$. Although all politicians know that low-ability sincere-myopic politicians flip-flop while low-ability opportunists and high-ability politicians maintain their initial positions, high-ability and low-ability politicians will again differ in their beliefs regarding their opponent's ability because private signals are correlated. Hence, a k -ability candidate's expectation of the citizen's posterior on ρ when she maintains her initial position is $(\hat{\beta}_k + (1 - \hat{\beta}_k)\theta)\Pr[\rho|r_{\bar{A}}, r_{\bar{A}}, l_\pi] + (1 - \hat{\beta}_k)(1 - \theta)\Pr[\rho|r_{\bar{A}}, r_L, l_\pi]$; her expectation of the citizen's posterior on λ when she flip-flops is $(\hat{\beta}_k + (1 - \hat{\beta}_k)\theta)\Pr[\lambda|r_L, r_{\bar{A}}, l_\pi] + (1 - \hat{\beta}_k)(1 - \theta)\Pr[\lambda|r_L, r_L, l_\pi]$.

Given these components of a candidate's expected approval rating, the incentive compatibility constraint which guarantees that a k -ability opportunist will maintain her initial position when it is common knowledge that all opportunists maintain their initial positions is:

$$\begin{aligned} & \alpha \left[(\hat{\beta}_k + (1 - \hat{\beta}_k)\theta)\Pr[\rho|r_{\bar{A}}, r_{\bar{A}}, l_\pi] + (1 - \hat{\beta}_k)(1 - \theta)\Pr[\rho|r_{\bar{A}}, r_L, l_\pi] \right] + (1 - \alpha)q_{\bar{A}} \\ & > \alpha \left[(\hat{\beta}_k + (1 - \hat{\beta}_k)\theta)\Pr[\lambda|r_L, r_{\bar{A}}, l_\pi] + (1 - \hat{\beta}_k)(1 - \theta)\Pr[\lambda|r_L, r_L, l_\pi] \right] + (1 - \alpha)q_L \text{ (B)} \end{aligned}$$

The LHS of (B) is a k -ability politician's expected approval rating from maintaining support for policy r . The RHS of (B) is a k -ability politician's expected approval rating from switching support to policy l . The proof of the second part follows trivially from two observations: (i) the LHS of (B) increases in $\hat{\beta}_k$ while the RHS of (B) decreases in $\hat{\beta}_k$, and (ii) $\hat{\beta}_H > \hat{\beta}_L$. Hence, if maintaining an initial position is optimal for the low-ability opportunist, then maintaining an initial position must also be optimal for a high-ability opportunist.

Lemma 1 tells us that both kinds of inefficiencies cannot arise in equilibrium. This leaves open the possibility for three combinations of opportunist behavior to constitute equilibria: (i) high-ability opportunists inefficiently flip-flop while low-ability

opportunists efficiently flip-flop, (ii) low-ability opportunists inefficiently maintain their initial positions while high-ability opportunists efficiently maintain their initial positions, and (iii) high-ability opportunists efficiently maintain their initial positions and low-ability opportunists efficiently flip-flop. The following propositions examine how the first two of these outcomes can constitute equilibria.

Proposition 1 establishes the circumstances under which high-ability opportunists can inefficiently flip-flop in equilibrium.

Proposition 1 (Opportunistic Policy Flip-flopping) *Provided conditions (E1) and (E2) are satisfied, and (K) holds for high-ability opportunists, there exists a pure-strategy equilibrium in which q_H -opportunists flip-flop inefficiently and q_L -opportunists flip-flop efficiently. Here flip-flopping signals to the citizen that a candidate is more likely to be an opportunist; maintaining an initial position identifies the citizen as high-ability.*

As mentioned earlier, conditions (E1) and (E2) establish informational efficient behavior. That is, they are incentive compatibility constraints. (E1) states that it is informationally efficient for a high-ability politician to maintain her initial position and therefore guarantees that a high-ability sincere-myopic politician will do so. (E2) states that it is informationally efficient for a low-ability politician to flip-flop and therefore guarantees that a low-ability sincere-myopic politician will do so. As mentioned above, condition (K) is the incentive compatibility constraint which guarantees that a k -ability opportunist will flip-flop. From Lemma 1, we know that when (K) holds for a high-ability opportunist, it must also hold for a low-ability opportunist. Hence, the low-ability opportunist's incentive compatibility constraint is slack. By construction, (E1), (E2), and (K) defined for a high-ability opportunist demarcate the region of the parameter space in which high-ability opportunists inefficiently flip-flop and low-ability opportunists efficiently flip-flop in equilibrium.

Here flip-flopping signals to the citizen that a candidate is more likely to be an opportunist because all opportunists flop-flop while only high-ability sincere-myopic candidates maintain their initial position. For the same reason, maintaining an initial position signals to the citizen that a candidate is of high ability. This captures the conventional wisdom view regarding policy flip-flops in the face of contradictory public opinion. First, a politician who maintains her initial position honestly believes in the policy she advocates. Second, a politician who flip-flops is likely to be doing so to improve her election prospects—she panders to public opinion.

Figure 3 illustrates Proposition 1. When an high-ability opportunist considers whether to flip-flop or maintain her initial position, she anticipates how the citizen will view her ex-post. If she maintains her initial position and her opponent flip-flops, the citizen’s ex-post view of her is point a ; if both she and her opponent maintain their initial positions, the citizen’s ex-post view of her is point b . The high-ability opportunist’s expectation of the citizen’s ex-post view of her when she maintains her initial position is point c . If she flip-flops and her opponent maintains her initial position, the citizen’s ex-post view of her is point d ; if both she and her opponent flip-flop, the citizen’s ex-post view of her is point e . The high-ability opportunist’s expectation of the citizen’s ex-post view of her when she flip-flops is point f .

A high-ability opportunist will flip-flop when point f yields a greater expected approval rating than point c . This situation occurs when the slope of the citizen’s approval-rate indifference curves is greater than the slope between points f and c . This intuition is the following. The slope between points f and c represents the trade-off a high-ability opportunist faces when she considers flip-flopping: by flip-flopping the candidate decreases the expected posterior the citizen will have regarding her ability to address the future issue, but increases the expected posterior the citizen will have regarding the appropriateness of her revised election-year position. When the citizen’s indifference curve is steeper than the line between f and c , the citizen

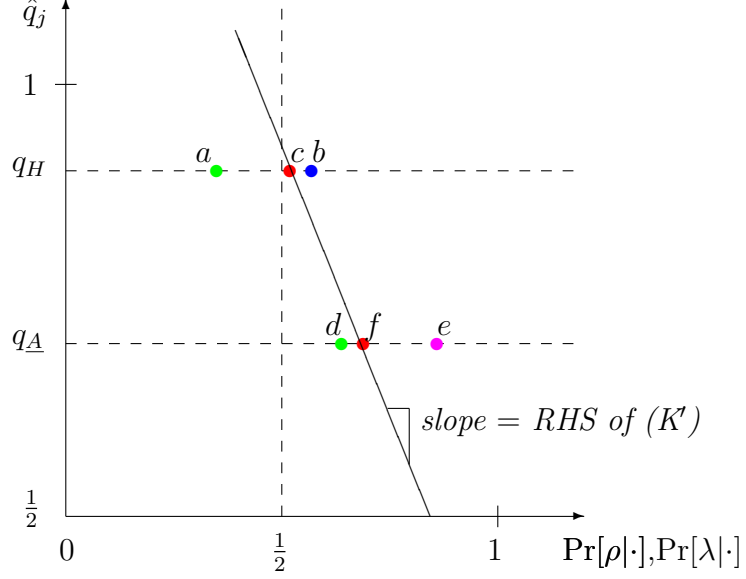


Figure 3: Opportunistic Policy Flop-flopping

cares little for the future issue, but cares greatly for the election-year issue. Hence, the high-ability opportunist will forgo the expected increase in her perceived ability to address the future issue appropriately by maintaining her initial position and instead flip-flop to obtain the expected increase in the citizen's belief that she has chosen the appropriate revised position on the election-year issue. This situation can be clearly seen if we rewrite condition (K) for the high-ability opportunist as:

$$\frac{\alpha}{1 - \alpha} > \frac{\frac{1-\beta}{1-\beta+\beta\theta}(q_H - q_L)}{(1 - \hat{\beta}_H + \hat{\beta}_H\theta) \left[\Pr[\lambda|q_A, q_A] - \Pr[\rho|q_H, q_A] \right] + \hat{\beta}_H(1 - \theta) \left[\Pr[\lambda|q_A, q_H] - \Pr[\rho|q_H, q_H] \right]} . \quad (K')$$

Condition (K') captures the intuition developed in Figure 3. The LHS of (K') is the citizen's marginal rate of substitution between having the election-year issue resolved appropriately and having the future issue resolved appropriately. The RHS of (K') is the slope between points f and c .

At this point, the reader may be curious if a pure-strategy equilibrium in

which all high-ability opportunists flip-flop can exist when the entire population of politicians consists of opportunists. The following corollary shows that this is indeed possible.

Corollary 1 *Even when $\theta = 1$, there exist circumstances in which all high-ability opportunists inefficiently flip-flop with full probability in equilibrium.*

To prove Corollary 1, evaluate (K') at $\theta = 1$. The RHS of the expression must be positive, otherwise we cannot rewrite (K) defined for a high-quality opportunist as (K'). It can be easily shown that the RHS is positive if the following condition holds.

$$\left(\frac{\pi}{1-\pi}\right)^2 > \frac{q_H}{1-q_H} \left(\frac{q_A}{1-q_A}\right)^3 \quad (3)$$

Condition 3 says that two public signals must be stronger than four private signals, one of high quality and three of “average” quality. By construction, when this condition is satisfied, the high-ability opportunist’s incentive compatibility constraint that guarantees inefficient flip-flopping will hold at $\theta = 1$.

We now examine inefficient equilibrium behavior by low-ability opportunists. Proposition 2 establishes that low-ability opportunists can inefficiently maintain their initial positions in equilibrium.

Proposition 2 (Opportunistic Policy Maintenance) *Provided conditions (E1) and (E2) are satisfied, and (B) holds for low-ability opportunists, there exists a pure-strategy equilibrium in which q_L -opportunists maintain policy inefficiently and q_H -opportunists maintain policy efficiently. Here flip-flopping signals to the citizen that a candidate is less likely to be an opportunist, but also identifies the candidate as low-ability.*

Conditions (E1) and (E2) again establish informational efficient behavior. As mentioned above, condition (B) is the incentive compatibility constraint which guarantees

that a k -ability opportunist will maintain her initial position. From Lemma 1, we know that when (B) holds for a low-ability opportunist, it must also hold for a high-ability opportunist. Hence the high-ability opportunist's incentive compatibility constraint is slack. By construction, (E1), (E2), and (B) defined for a low-ability opportunist demarcate the region of the parameter space in which high-ability opportunists efficiently maintain their initial positions and low-ability opportunists inefficiently maintain their initial positions in equilibrium.

Here maintaining the initial policy position signals to the citizen that a candidate is more likely to be an opportunist because all opportunists maintain their initial positions while only low-ability sincere-myopic candidates flip-flop. For the same reason, flip-flopping signals to the citizen that a candidate is of low-ability. This captures the Majumdar-Mukand view regarding policy flip-flops in the face of contradictory information public opinion. First, politicians who flip-flop have learned that their initial position was mistaken and honestly believe in the new policy they advocate. Second, politicians who maintain their initial positions are likely to be doing so to improve their election prospects—they fear that they will be perceived as being of incompetent if they flip-flop.

Figure 4 illustrates Proposition 2. When an low-ability opportunist considers whether to flip-flop or maintain her initial position, she anticipates how the citizen will view her ex-post. If she maintains her initial position and her opponent flip-flops, the citizen's ex-post view of her is point a' ; if both she and her opponent maintain their initial positions, the citizen's ex-post view of her is point b' . The low-ability opportunist's expectation of the citizen's ex-post view of her when she maintains her initial position is point c' . If she flip-flops and her opponent maintains her initial position, the citizen's ex-post view of her is point d' ; if both she and her opponent flip-flop, the citizen's ex-post view of her is point e' . The high-ability opportunist's

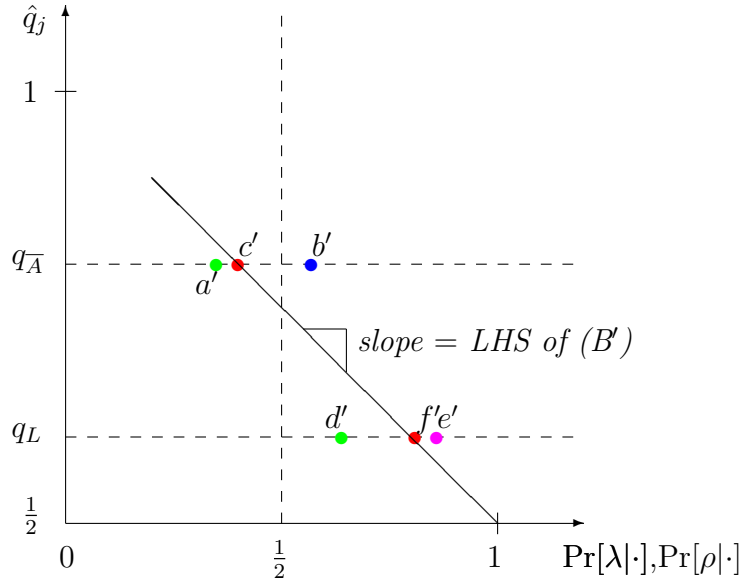


Figure 4: Opportunistic Policy Maintenance

expectation of the citizen's ex-post view of her when she flip-flops is point f' .

A low-ability opportunist will maintain her initial position when point c' yields a greater expected approval rating than point f' . This situation occurs when the slope of the citizen's approval-rate indifference curves are less than the slope between points c' and f' . This intuition is similar to Proposition 1. The slope between points c' and f' represents the trade-off a low-ability opportunist faces when she considers maintaining her initial position: by maintaining her initial position the candidate decreases the expected posterior the citizen will have regarding the appropriateness of her revised election-year position, but increases the expected posterior the citizen will have regarding her ability to appropriately address the future issue. When the citizen's indifference curve is flatter than the line between c' and f' , the citizen cares little for the election-year issue, but cares greatly for the future issue. Hence, the low-ability opportunist will forgo the expected increase in the citizen's belief that she has chosen the appropriate revised position on the election-year issue that comes from flip-flopping and instead maintains her initial position to obtain the expected increase in her perceived ability to address the future issue. This situation can be

clearly seen if we rewrite condition (B) for the low-ability opportunist as:

$$\frac{\frac{\beta}{(1-\beta)\theta+\beta}(q_H - q_L)}{(\hat{\beta}_L + (1 - \hat{\beta}_L)\theta) \left[\Pr[\lambda|q_L, q_{\bar{A}}] - \Pr[\rho|q_{\bar{A}}, q_{\bar{A}}] \right] + (1 - \hat{\beta}_L)(1 - \theta) \left[\Pr[\lambda|q_L, q_L] - \Pr[\rho|q_{\bar{A}}, q_L] \right]} > \frac{\alpha}{1 - \alpha} \quad (\text{B}')$$

Condition (B') captures the intuition developed in Figure 4. The LHS of (B') is the slope between points f' and c' . The RHS of (B') is the citizen's marginal rate of substitution between having the election-year issue resolved appropriately and having the future issue resolved appropriately.

At this point, a natural question to ask is whether the equilibria described in Propositions 1 and 2 are disjoint. The following corollary and conjecture address this issue.

Corollary 2 *These equilibria are disjoint when $\theta = 0$.*

To prove Corollary 2, we need to show that the boundaries of these equilibria do not overlap at $\theta = 0$. The RHS of (K') is the boundary of the equilibrium in Proposition 1. The LHS of (B') is the boundary of the equilibrium in Proposition 2. Therefore, if the RHS of (K') is greater than the LHS of (B') at $\theta = 0$, then these equilibria are disjoint. When evaluating the RHS of (K') at the $\theta = 0$, notice that $q_{\bar{A}}$ is equal to q_L . Similarly, when evaluating the LHS of (B') at $\theta = 0$, notice that $q_{\bar{A}}$ is equal to q_H . It follows that the RHS of (K') is greater than the LHS of (B') because $1 > q_H > q_L > \frac{1}{2}$.

Although we can easily show that these equilibria are disjoint at $\theta = 0$, it is difficult to do so for all values of $\theta \in [0, 1]$. Instead of proving this claim, we make the conjecture that these equilibria are disjoint for all $\theta \in [0, 1]$ and provide the intuition for this claim along with two numerical examples.

Conjecture 1 *These equilibria are disjoint for all values of $\theta \in [0, 1]$.*

From Corollary 2, we know that the equilibria in Propositions 1 and 2 are disjoint at $\theta = 0$. To understand the intuition for Conjecture 1, we examine how an opportunist's trade-off between flip-flopping and maintaining her initial position change as θ increases (i.e., as the population of opportunists increases). We first consider the equilibrium described in Proposition 1 where all high-ability opportunists inefficiently flip-flop and then consider the equilibrium described in Proposition 2 where all low-ability opportunists inefficiently maintain their initial positions.

Suppose all opportunists flip-flop in equilibrium. Here a high-ability opportunist forgoes the opportunity to identify herself as a high-ability politician by maintaining policy r and instead chooses to flip-flop to the policy l which the citizen believes is more appropriate. As θ increases, the citizen's posterior that l is appropriate falls because he knows that more high-ability opportunists are flip-flopping to l when it is informational efficient to maintain policy r . Therefore, roughly speaking, the high-ability opportunist has less of an incentive to flip-flop. That is, the region of the parameter space that supports an equilibrium in which all high-ability opportunists flip-flop shrinks as θ increases.

Now suppose all opportunists maintain their initial positions in equilibrium. Here a low-ability opportunist forgoes the opportunity to choose the policy the citizen believes is more appropriate by flip-flopping to l and instead chooses to increase the citizen's perception of her ability by maintaining policy r and pooling with high-ability politicians. As θ increases, the citizen's posterior regarding the ability of a candidate who maintains policy r falls because he knows that more low-ability opportunists are maintaining policy r and pooling with high-ability politicians. Therefore, roughly speaking, the low-ability opportunist has less of an incentive to flip-flop. That is, the region of the parameter space that supports an equilibrium in which all low-ability opportunist maintain their initial positions shrinks as θ increases.

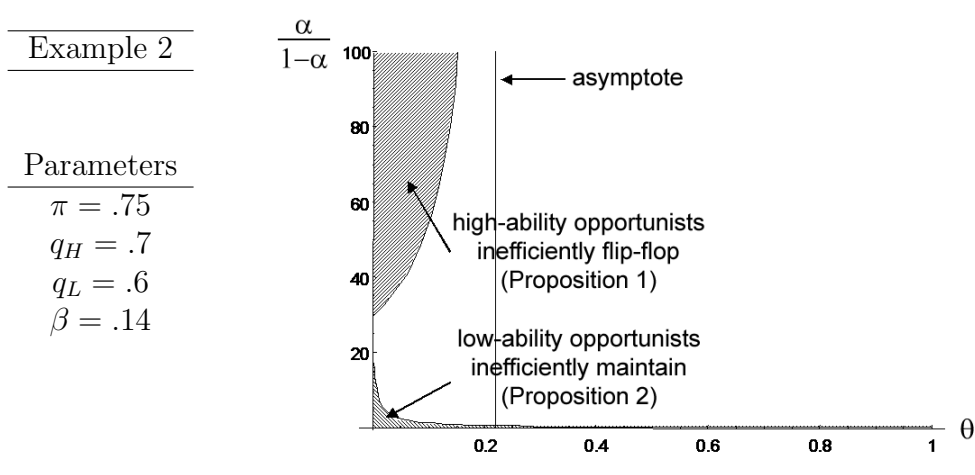
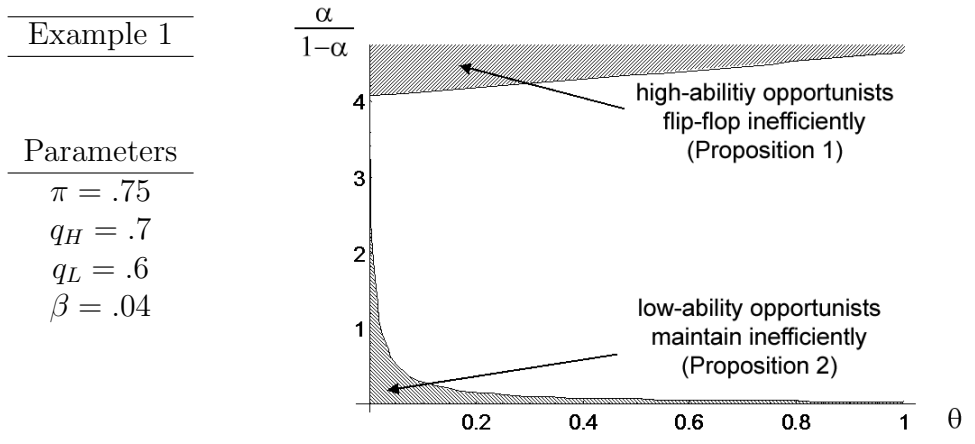


Figure 5: Two Numerical Examples

The rough analysis above suggests that the equilibria in Propositions 1 and 2 are disjoint. Further support for Conjecture 1 comes from numerical examples; we provide two in Figure 5. The two examples differ only in the population of high-ability politicians. Example 1 is a setting where condition (3) is satisfied; example 2 is a setting where condition (3) is not satisfied. Notice that when condition (3) fails, there is an asymptote which bounds the region in which all opportunists flip-flop with full probability. In both examples, the boundary of the region supporting the pure-strategy equilibrium where all opportunists flip-flop is increasing and the boundary of

the region supporting the pure-strategy equilibrium where all opportunists maintain their initial positions is decreasing. Therefore, the parameters that support the equilibrium of Proposition 1 and 2 are disjoint. This reflects the rough intuition discussed above.

The reader may be curious about the equilibria that exist in the unshaded regions of Examples 1 and 2. Much of this region supports mixed-strategy equilibria in which either the high-ability opportunist behaves inefficiently with some positive probability or the low-ability opportunist behaves inefficiently with some positive probability. To understand the intuition for existence of this kind of mixed-strategy equilibria, consider the following. Suppose α is such that there exists a pure-strategy equilibrium in which all high-ability opportunist inefficiently flip-flop when $\theta = 0$, but that no such equilibrium exists when $\theta = 1$. As θ increase from 0, we will eventually hit the boundary of the region that supports this pure-strategy equilibrium. As θ increase beyond this point, there always exists a mixed-strategy equilibrium in which high-ability opportunists randomize between flip-flopping to l and maintaining r in such a way as to stay on the boundary. Similar logic holds for the existence of equilibria in which low-ability opportunists mix.

2.4 Conclusion

We have presented a model which unifies two divergent views on policy flip-flopping. The conventional wisdom view says that politicians who flip-flop are more likely to be opportunists. We have shown that this occurs when the citizen cares relatively more about how the election-year issue is addressed. The Majumdar-Mukand view says that politicians who flip-flop are of low-ability. Therefore, opportunists would never flip-flop for fear of looking incompetent. We have shown that this occurs when

the citizen cares relatively more about how the future issue is addressed.

References

Majumdar, Sumon and Sharun Mukand. “Policy Gambles.” *American Economic Review* 94 (2004): 1207 – 1222.

Polak, Benjamin and Adam Brandenburger. “When Managers Cover Their Posteriors: Making the Decisions the Market Wants to See.” *Rand Journal of Economics* 27 (1996): 523-541.

Chapter 3

Strategic Export Promotion

in the Presence of Special Interests

Keywords: strategic trade, export subsidies, political agency, corruption, elections, signaling

JEL classification: D72, D73, D82, F13

3.1 Introduction

The desirability of free trade is one of the profession's most cherished beliefs. In most situations of perfect competition, the optimal policy for governments is free trade. However, Brander and Spencer (1985) shows that in situations of imperfect competition, the optimal policy for governments may be an export subsidy. Such strategic trade policies are not without their problems. Eaton and Grossman (1986) shows that the nature of the policy is quite sensitive to the assumption regarding how firms compete. If firms compete in quantities, then the optimal policy is generally a subsidy. However, if firms compete in prices, then the optimal policy is generally a tax. In this manner, the knowledge required for a government implementing strategic trade policy is formidable.

Even if governments know how firms compete in an industry in which strategic trade policy is being considered, there is a strong political economy argument against the implementation of such a policy. Grossman criticizes strategic trade policy saying:

Experience has shown that the trade policy apparatus is susceptible to the political pressures of special interest groups, and that the outcomes often fail to take adequate account of the interests of consumers... The risk that any scheme of targeted export promotion would fall prey to much the same sort of special interest pressures is cause for grave concern. If an apparatus for discretionary industrial policy of this type were to be erected, each and every export sector would have ample incentive to argue the (alleged) merits of its own case for subsidization. And even if the policy analyst could somehow solve the difficult technical problems of identifying industries worthy of promotion, there could be little guarantee that these would be the ones to emerge from a politically influenced process of selection. (Grossman 1986, p.65)

To our knowledge, Grossman's political economy critique of sophisticated export subsidization has not been formally modeled. In doing so, we draw on a political economy model developed in Coate and Morris (1995) which analyzes transfers from politicians to special interest groups. In particular, they study how public projects

may be undertaken by politicians to benefit a special interest at the expense of the citizenry even though it would be more efficient to transfer income directly from the citizenry to the special interest. They consider a model in which an incumbent politician is seeking reelection. The incumbent may be good in which case she acts to benefit the citizenry or she may be bad in which case she acts to benefit a special interest. The incumbent politician has a discrete decision regarding whether or not to implement a project and a continuous decision regarding the level of a non-negative direct transfer from the citizenry to a special interest. The citizenry observe the incumbent's decisions and the results of the project if it is implemented. They elect the candidate, the incumbent or a randomly chosen challenger, that they believe is more likely to be good. To analyze the use of the project by the bad incumbent as an inefficient means for delivering income to the special interest, two instruments are available to the incumbent: the project and the direct transfer.

Our model differs in some notable aspects. We study how politicians use export subsidies as a means for delivering income to a special interests. We consider a similar model in which an incumbent politician, who is either good or bad, is seeking reelection. The incumbent, however, has two continuous decisions: the level of a non-negative subsidy to a domestic firm and the level of a non-negative transfer to a special interest. The citizenry observe only the incumbent's decisions and elect the candidate, the incumbent or a randomly chosen challenger, that they believe is more likely to be good. We include the direct transfer instrument so that when good and bad incumbents separate themselves in equilibrium, bad incumbent does so efficiently.

When the citizenry are uncertain about whether the incumbent is good or bad and uncertain about whether or not subsidies are in their benefit, then separating and partial-pooling equilibria arise. We find that if the good incumbent can separate herself from the bad incumbent she will do so by selecting a sub-optimally low subsidy

for the domestic firm. Furthermore, we find that when she is unable to separate herself, then the bad incumbent will, in certain circumstances, use the subsidy to deliver income to the special interest at the expense of the citizenry. This is the danger that regimes of export subsidization present to the citizenry.

This paper is organized as follows. Section 2 outlines the model. In section 3, we analyze the separating and partial-pooling equilibria of the model and discuss the welfare implications of export subsidies. In section 4, we extend the model to a setting with a continuum of industries each with the potential of receiving export subsidies. Section 5 concludes.

3.2 The Model

We construct a two-country model in which two firms, one domestic and one foreign, are Cournot competitors and only operate in an external third market.¹ There are two periods. In the first period, an incumbent politician leads the domestic government and must decide how much to subsidize the domestic firm. Subsidizing the domestic firm may or may not be of benefit to the domestic citizen, but it is always of some benefit to a special interest group. In addition, the incumbent must also decide whether or not to give a direct transfer to the special interest. After the subsidy and transfer decisions have been made, an election is held between the incumbent and a randomly chosen challenger. The domestic citizen alone possesses the power to vote and determines the outcome of the election. In the second period, the winner of the election chooses a new subsidy level for the domestic firm and decides whether or not to implement a second direct transfer to the special interest.

The Domestic Citizen and the Special Interest. Suppose there are two coun-

¹This is the setting presented in Brander and Spencer (1985) and is most conducive to the use of export subsidies.

tries, one domestic and one foreign, each with a corresponding firm. The two firms are Cournot competitors and only operate in a third market that is external to the two countries. The domestic firm may or may not receive a non-negative subsidy from the domestic government. The foreign firm receives no such subsidy from the foreign government. The domestic firm and foreign firm profit functions are respectively,

$$\pi(q, q_f, s) = p(Q)q - cq + sq \quad \text{and} \quad (1)$$

$$\pi_f(q, q_f, s) = p(Q)q_f - c_f q_f, \quad (2)$$

where all f -subscript terms denote foreign firm variables, q is the firm's production, $Q = q + q_f$ is the sum of domestic and foreign production, $p(Q)$ is the inverse demand function, c is the firm's marginal cost, and s is the subsidy. We assume that q and q_f are strategic substitutes:

$$\begin{aligned} \frac{\partial^2 \pi}{\partial q \partial q_f} &= p''q + p' < 0 \\ \frac{\partial^2 \pi_f}{\partial q \partial q_f} &= p''q_f + p' < 0 \end{aligned}$$

This assumption guarantees that each firm's best response function is downward sloping so that the optimal strategic trade policy in a Cournot duopoly is generally a subsidy. Furthermore, this assumption is a sufficient condition for each firm's objective function to be concave so that the first order condition yields a maximum. Notice that the strategic substitutes assumption can be violated if the inverse demand function $p(Q)$ is too convex.

The domestic citizen owns a share $\alpha \in (0, 1)$ of the domestic firm and is therefore entitled to an α -share of the domestic firm's profits. The remaining $(1 - \alpha)$ -share of the domestic firm's profits go to the special interest. The cost of any subsidy

is borne fully by the domestic citizen. Furthermore, the domestic government has the ability to transfer a non-negative amount of income directly from the domestic citizen to the special interest. Hence, the per period utility functions of the domestic citizen and special interest are respectively,

$$u_c = \alpha\pi - sq - T \quad \text{and} \quad (3)$$

$$u_{si} = (1 - \alpha)\pi + T, \quad (4)$$

where T is the direct transfer.

After the domestic government's first-period subsidy and transfer decisions, but prior to the distribution of profits from the domestic firm, the domestic citizen makes his only decision in the game: whether or not to reelect the incumbent politician who leads the domestic government. The special interest has no decisions to make.

Government Policies and Policy Uncertainty. In each period, the politician who leads the government must decide whether or not to implement a non-negative direct transfer to the special interest. A direct transfer takes T from the domestic citizens and gives it to the special interest. Hence, if implemented, a direct transfer clearly benefits the special interest at the expense of the domestic citizens.

In addition to her transfer decision(s), in the each period the politician who leads the government must also choose the level of the subsidy that the domestic firm receives. Subsidizing the firm is always beneficial to the special interest. To see this, differentiate equation (4) with respect to s . Then using properties of the Cournot equilibrium obtain:

$$\frac{du_{si}}{ds} = (1 - \alpha)p'q^* \frac{dq_f^*}{ds} + (1 - \alpha)q^*, \quad (5)$$

where all starred variables denote equilibrium values. Notice that $p' < 0$ and with strategic substitutes $dq_f^*/ds < 0$. Hence, this expression is unambiguously positive and establishes our first lemma.

Lemma 1 *Subsidizing the firm is always beneficial to the special interest.*

Unlike the special interest, the domestic citizen may or may not benefit from subsidizing the firm—this depends on α and the inverse demand function $p(Q)$. To see this, obtain the first order condition of (3) with respect to s . Then, using properties of the Cournot equilibrium, express the optimal subsidy from the domestic citizen's point of view as:

$$s^* = \left(\frac{dq^*}{ds} \right)^{-1} \left(\alpha \left[p' q^* \frac{dq_f^*}{ds} \right] + (\alpha - 1) q^* \right) \quad (6)$$

Note that with strategic substitutes $dq^*/ds > 0$. Hence, the coefficient is positive. Furthermore, notice again that $p' < 0$ and with strategic substitutes $dq_f^*/ds < 0$. Hence, the first term is positive. The second term, however, is negative because $\alpha \in (0, 1)$. Therefore, if we were to ignore the non-negative restriction on the subsidy, the sign of the optimal subsidy from the domestic citizens' point of view is ambiguous. This leads to our second lemma.

Lemma 2 *Fix the marginal costs of the domestic and foreign firms, c and c_f . Then there exists an $\underline{\alpha} \in (\frac{1}{2}, \frac{3}{4}]$ and $\bar{\alpha} \in (\frac{3}{4}, 1]$ with the following properties:*

1. *If $\alpha \geq \bar{\alpha}$, then the domestic citizen always benefits from a positive subsidy.*
2. *If $\alpha \leq \underline{\alpha}$, then the domestic citizen never benefits from a positive subsidy.*
3. *If $\underline{\alpha} < \alpha < \bar{\alpha}$, then there exist two inverse demand functions $\underline{p}(Q)$ and $\bar{p}(Q)$ such that the optimal subsidy from the domestic citizen's point of view is $s = 0$ under $\underline{p}(Q)$ and $s > 0$ under $\bar{p}(Q)$.*

We assume $\alpha \in (\underline{\alpha}, \bar{\alpha})$ so that the domestic citizen's preference regarding a subsidy will depend his share of the domestic firm's profits and on the inverse demand function. In our model, the inverse demand function is stochastic. In one state it is $\bar{p}(Q)$ and hence a subsidy benefits the domestic citizen; this state occurs with probability ϕ^s . In the other state it is $\underline{p}(Q)$ and hence a subsidy harms the domestic citizen; this state occurs with probability $(1 - \phi^s)$.

The incumbent and the special interest observe the state, i.e., the realization of the inverse demand function. The domestic citizen, however, does not discover the realization of the inverse demand function until after the election is held when profits are distributed. Therefore, at the time of the election the domestic citizen does not know if a subsidy is beneficial. In this manner, the domestic citizen faces what Coate and Morris (1995) refers to as “policy uncertainty.”

Politicians and Politician Uncertainty. Politicians can be one of two possible types: *good* or *bad*. A good politician who holds office cares only about the welfare of the domestic citizen and hence her per period utility is $u_g(s, T, p(Q)) = u_c(s, T, p(Q))$. A bad politician who holds office cares only about the welfare of the special interest and hence her per period utility is $u_b(s, T, p(Q)) = u_{si}(s, T, p(Q))$. Politicians receive zero utility when they are out of office regardless of their type. In addition, politicians discount future payoffs using discount factor δ .

The domestic citizen faces what Coate and Morris (1995) calls “politician uncertainty.” That is, he does not observe politicians' types. Instead he has an initial belief regarding a politician's type—this is also called a politician's initial reputation. Specifically, the domestic citizen believes the incumbent is good with probability $\theta_I \in (0, 1)$. Similarly, he believes the challenger is good with probability $\theta_C \in (0, 1)$. The incumbent knows her initial reputation. After the incumbent makes her first period subsidy and direct transfer decisions, a challenger is chosen, i.e., θ_C is drawn

from some smooth and increasing cumulative distribution function $\lambda(\theta)$ with $\lambda(0) = 0$.

The Game. This model describes a two-period game between the incumbent, the challenger, and the domestic citizen. At the beginning of the first period, nature chooses the inverse demand function $p(Q) \in \{\underline{p}(Q), \bar{p}(Q)\}$ and the incumbent's type, *good* or *bad*. The incumbent observes both of nature's choices. The domestic citizen has belief θ_I that the incumbent is good and only knows that $p(Q) = \bar{p}(Q)$ with probability ϕ^s . After nature has moved, the incumbent must choose the subsidy level for the domestic firm and whether or not to implement the direct transfer. The domestic citizen observes the incumbent's decisions. In keeping with Coate and Morris (1995), we refer to (s_1, T_1) as the incumbent's first-period record where $s_1 \in [0, \infty)$ is the incumbent's first-period subsidy decision and $T_1 \in [0, \infty)$ is the incumbent's first-period direct transfer decision.

After the domestic citizen observes the incumbent's first-period record, nature randomly chooses the challenger and an election is held. As mentioned above, the challenger has an initial reputation θ_C which is drawn from the cumulative distribution function $\lambda(\theta)$. The domestic citizen bases his voting decision on the incumbent's initial reputation and first-period record and the challenger's initial reputation. At the end of the first period, the domestic and foreign firms compete and profits are distributed (i.e., the domestic citizen discovers $p(Q)$). In the second period, the elected politician chooses a new subsidy level for the domestic firm $s_2 \in [0, \infty)$ and decides whether or not to implement a second direct transfer $T_2 \in [0, \infty)$. Then the game ends.

In each period, the politician who leads the government faces the following constraint: her subsidy and transfer decisions cannot bankrupt the domestic citizen. That is, the subsidy and transfer chosen by the politician in each period must yield

the domestic citizen non-negative utility. Formally,

$$u_c(s_\tau, T_\tau, p) = \alpha \pi(s_\tau, p) - s q(s_\tau, p) - T_\tau \geq 0 \quad \forall \tau \in \{1, 2\} \text{ and } \forall p \in \{\underline{p}, \bar{p}\} \quad (7)$$

In this model, a Perfect Bayesian Nash Equilibrium (PBNE) consists of: (1) the domestic citizen's beliefs, (2) a strategy for the domestic citizen, (3) a strategy for the incumbent, and (4) a strategy for the challenger, which satisfy the following properties. First, the domestic citizen's beliefs must be consistent with the incumbent's strategy—they must be formed using Bayes Rule where possible. Second, the domestic citizen's strategy must be optimal given his beliefs and the strategies of the incumbent and the challenger. Third, the incumbent's strategy must be optimal given the beliefs and strategy of the domestic citizen and the strategy of the challenger. Fourth, the challenger's strategy must be optimal.

The domestic citizen's beliefs specify the probabilities regarding the realizations of $p(Q)$ and the incumbent's type. The belief as to the realization of $p(Q)$ is based solely on the ϕ^s . The belief as to the realization of the incumbent's type is based on the incumbent's initial reputation and her first-period record.

A strategy for the domestic citizen specifies the probability with which to vote for the incumbent. This probability depends on the incumbent's initial reputation and her first-period record.

A strategy for the incumbent is a rule that specifies an action or actions to be taken in the first period and in the second period if reelected. In the each period, the strategy specifies a subsidy level and whether or not to implement a direct transfer given the incumbent's type and the realization of $p(Q)$.

A strategy for the challenger is a rule that specifies, if elected, a subsidy level and whether or not to implement a direct transfer given the challenger's type and the

realization of $p(Q)$.

3.3 Equilibrium Subsidies and Direct Transfers

Using backward induction, we show that there are two kinds of pure-strategy PBNE in this game. First, separating equilibria exist with the following first-period behavior. The good incumbent never implements the direct transfer and chooses a suboptimal subsidy level when the inverse demand function is $p(Q) = \bar{p}(Q)$ and the optimal subsidy when the inverse demand function is $p(Q) = \underline{p}(Q)$. The bad incumbent behaves to maximize her first-period payoff which involves the use of a direct transfer and subsidy. Second, partial-pooling equilibria exist with the following first-period behavior. The good incumbent chooses the optimal subsidy from the domestic citizen's point of view given the realization of the inverse demand function $p(Q)$, while the bad incumbent always chooses the subsidy which the good incumbent chooses when inverse demand function is $p(Q) = \bar{p}(Q)$. Neither type of incumbent ever implements the first-period direct transfer.

Politicians' Second-Period Behavior. Regardless of whether the incumbent or the challenger is elected, the politician who leads the government during the second period chooses the action that maximizes her one period utility given her type and the realization of $p(Q)$.

A good politician cannot increase her one period utility by implementing a positive direct transfer regardless of the inverse demand function's realization. Hence, she will select a zero direct transfer $T_g(p) = 0$ for all inverse demand functions. She maximizes her one period utility by choosing subsidy level $s_g^*(p)$ which is defined as

$$s_g^*(p) = \operatorname{argmax}_{s \in [0, \infty)} u_g(s, 0, p) \quad . \quad (8)$$

Let $u_g^*(p)$ denote the resulting utility level. From the way in which the inverse demand function is defined, we know that $s_g^*(\underline{p}) = 0$ and $s_g^*(\bar{p}) > 0$. This strategy does not violate the no-bankruptcy condition because the good politician's per period utility function is identical to that of the domestic citizen.

A bad politician maximizes her one period utility by using both the direct transfer and the subsidy. First, notice that $du_b/ds < 1$ for all α and $p(Q)$ which satisfy point 3 of Proposition 2 and that $du_b/dT = 1$. Therefore, the bad politician will exhaust the direct transfer before choosing the subsidy level. Due to the no-bankruptcy condition, the greatest direct transfer is $T_b(p) = \alpha \pi(s, p) - s q(s, p)$ which leaves the domestic citizen with zero one period utility regardless of the inverse demand function's realization. Given this direct transfer, the bad politician maximizes her one period utility by choosing subsidy level $s_b^*(p)$ which is defined as

$$s_b^*(p) = \operatorname{argmax}_{s \in [0, \infty)} u_b(s, T_b(p), p) \quad . \quad (9)$$

Let $u_b^*(p)$ denote the resulting utility level. Notice that $s_b^*(p)$ is the "standard" strategic trade subsidy found in Brander and Spencer (1985) given the inverse demand function is $p(Q)$.

The Domestic Citizen's Behavior. By backward induction, the domestic citizen's knows that he is better off in the second period if he elects a good politician than if he elects a bad politician. Therefore, he elects the politician he believes is more likely to be good. Given an incumbent's first-period record (s_1, T_1) , the domestic citizen estimates the probability that the incumbent is good, $\beta(s_1, T_1)$. As mentioned above, he is also aware of the challenger's initial reputation θ_C . Hence, the domestic citizen will reelect the incumbent if and only if $\beta(s, T_1) > \theta_C$. The probability that the incumbent is reelected is given by $\lambda(\beta(s_1, T_1))$ because the challenger's initial

reputation is randomly drawn from cumulative distribution function $\lambda(\theta)$.

The Incumbent's First-Period Behavior and the Beliefs of the Domestic

Citizen. Consider the incumbent's payoff when the inverse demand function is $p(Q)$.

If she is good, then her payoff as a function of her first-period subsidy and transfer decisions is

$$U_g(s_1, T_1, p) = u_g(s_1, T_1, p) + \delta \lambda(\beta(s_1, T_1)) u_g(s_g^*(p), 0, p). \quad (10)$$

If she is bad, then her payoff as a function of her first-period subsidy and transfer decisions is

$$U_b(s_1, T_1, p) = u_b(s_1, T_1, p) + \delta \lambda(\beta(s_1, T_1)) u_b(s_b^*(p), T_b(p), p). \quad (11)$$

Clearly, the optimal strategy for the incumbent will depend on the beliefs of the domestic citizen. In a PBNE, for any first-period record that occurs along the equilibrium path, these beliefs must be formed from the incumbent's strategy using Bayes's rule. However, for first-period records which do not occur along the equilibrium path, these beliefs are not restricted. This gives rise to multiple equilibria, many of which consist of unusual behavior on the part of the incumbent such as the good incumbent's implementing a positive first-period direct transfer. In order to rule out these bizarre equilibria, we assume that first-period records with a greater direct transfer cannot result in a greater belief that the incumbent is good. This is a monotonic beliefs property that is similar to the one found in Coate and Morris (1995). Formally, we assume that for any pair of first-period records (s, T) and (s, T') such that $T' > T$, then $\beta(s, T') < \beta(s, T)$.

In any equilibrium satisfying this monotonic beliefs property, the good incum-

bent will never implement a first-period direct transfer because doing so will lower her first-period utility and lower the probability that she is reelected, $\lambda(\beta(s_1, T_1))$. Therefore, the good incumbent's first-period record must be of the form $(s, 0)$ where $s \in [0, \infty)$ satisfies the no-bankruptcy condition. It then follows that if the domestic citizen observes a positive first-period direct transfer, he will believe that the incumbent is bad with probability one and vote her out of office. Hence, if a bad incumbent ever implements a positive first-period direct transfer, she will be acting to maximize her one period utility because she knows that she will be voted out of office with certainty. Therefore, the bad incumbent's first-period record must be of the form $(s_b^*(p), T_b(p))$ or $(s, 0)$ where $s \in [0, \infty)$ satisfies the no-bankruptcy condition.

Let $\sigma_g(s, T, p)$ be the probability that the good incumbent chooses the first-period record (s, T) when the inverse demand function is $p(Q)$. Similarly, let $\sigma_b(s, T, p)$ be the probability that the bad incumbent chooses the first-period record (s, T) when the inverse demand function is $p(Q)$. Given this notation regarding the incumbent's first-period behavior, we can express the domestic citizen's equilibrium path beliefs using Bayes's Rule. From the above discussion, we know that first-period records that appear on the equilibrium path must be of the form $(s, 0)$ or (s, T_1) where $s \in [0, \infty)$ and $T_1 > 0$. If the domestic citizen observes a first-period record of the form $(s, 0)$ in equilibrium, he forms the belief

$$\beta(s, 0) = \frac{\theta_I [\phi^s \sigma_g(s, 0, \bar{p}) + (1 - \phi^s) \sigma_g(s, 0, \underline{p})]}{\theta_I [\phi^s \sigma_g(s, 0, \bar{p}) + (1 - \phi^s) \sigma_g(s, 0, \underline{p})] + (1 - \theta_I) [\phi^s \sigma_b(s, 0, \bar{p}) + (1 - \phi^s) \sigma_b(s, 0, \underline{p})]}.$$

If the domestic citizen observes a first-period record of the form (s, T_1) with $T_1 > 0$ in equilibrium, he forms the belief $\beta(s, T_1) = 0$.

The last task involved in finding the PBNE of this game is to solve for the values of $\sigma_g(s, T, p)$ and $\sigma_b(s, T, p)$ which maximize the incumbent's payoff given the beliefs they imply. In signaling games, there is generally a multiplicity of equilibria.

That is, there are many possible values of $\sigma_g(s, T, p)$ and $\sigma_b(s, T, p)$ with consistent beliefs which maximize the incumbent's payoff. Our game is no different, except that we restrict attention to the incumbent's pure strategies. We now analyze the domestic citizen's beliefs and incumbent's first-period behavior which generate separating equilibria and later analyze the domestic citizen's beliefs and incumbent's first-period behavior which generate partial-pooling equilibria.

3.3.1 Separating Equilibria

Our first proposition expresses the potential for a separating equilibrium: the good incumbent can separate herself from the bad incumbent by choosing a suboptimal subsidy when the inverse demand function is $\bar{p}(Q)$ and the optimal subsidy when it is $\underline{p}(Q)$.

Before discussing the separating equilibrium, we make the following two assumptions concerning the bad incumbent's preferences:

$$u_b^*(p) \geq u_b(s_g^*(\underline{p}), 0, p) + \delta u_b^*(p) \quad \forall p \in \{\underline{p}, \bar{p}\} \quad (\text{A1})$$

$$u_b^*(p) \leq u_b(s_g^*(\bar{p}), 0, p) + \delta u_b^*(p) \quad \forall p \in \{\underline{p}, \bar{p}\} \quad (\text{A2})$$

The first assumption is a necessary condition for a separating equilibrium. It states that the first-period record $(s_g^*(\underline{p}), 0)$ may separate the good and bad incumbents because the bad incumbent prefers the first-period record which maximizes her one period utility to the first-period record $(s_g^*(\bar{p}), 0)$, regardless of the inverse demand function's realization. In words this assumption says that if the domestic citizen were to reelect the incumbent with probability one upon seeing no first-period direct transfers, then the bad incumbent prefers the first-period record which maximizes her one period utility to the first-period record chosen by the good incumbent

under $p(Q) = \underline{p}(Q)$.

The second assumption is a non-triviality assumption. It states that the first-period record $(s_g^*(\bar{p}), 0)$ does not separate the good and bad incumbents because the bad incumbent prefers this first-period record to the one which maximizes her one period utility, regardless of the inverse demand function's realization. In words this assumption says that if the domestic citizen were to reelect the incumbent with probability one upon seeing no first-period direct transfers, then bad incumbent prefers the first-period record chosen by the good incumbent under $p(Q) = \bar{p}(Q)$ to the first-period record which maximizes her one period utility.

Under assumptions (A1) and (A2), there exists an \hat{s} with $0 < \hat{s} < s_g^*(\bar{p})$ defined as

$$\begin{aligned} \hat{s} = \min\{x, y\} \quad \text{where} \quad x \quad & \text{s.t.} \quad u_b^*(\underline{p}) = u_b(x, 0, \underline{p}) + \delta u_b^*(\underline{p}) \\ y \quad & \text{s.t.} \quad u_b^*(\bar{p}) = u_b(y, 0, \bar{p}) + \delta u_b^*(\bar{p}). \end{aligned}$$

The values x and y exist and lie in the interval $[0, s_g^*(\bar{p})]$ because the expression $u_b(s, 0, p) + \delta u_b^*(p)$ is continuous and strictly increasing in s regardless of the inverse demand function's realization. We now state our first proposition.

Proposition 1 (Suboptimal Use of Export Subsidy) *Under assumptions (A1) and (A2), there exists a separating equilibrium in which the bad incumbent chooses the first-period record which maximizes her one period utility while the good incumbent chooses the first-period record $(s_g^*(\underline{p}), 0)$ when $p(Q) = \underline{p}(Q)$ and the first-period record $(\hat{s}, 0)$ when $p(Q) = \bar{p}(Q)$, provided the following condition is satisfied:*

$$u_g^*(\bar{p}) \leq u_g(\hat{s}, 0, \bar{p}) + \delta u_g^*(\bar{p}) \quad (\text{T1})$$

Moreover, this is the only separating equilibrium that satisfies the Intuitive Criterion.

Proof. Suppose in equilibrium the good incumbent chooses $(\hat{s}, 0)$ when $p(Q) = \bar{p}(Q)$ and $(s_g^*(\underline{p}), 0)$ when $p(Q) = \underline{p}(Q)$ and that bad incumbent chooses $(s_b^*(p), T_b(p))$ in both states. This yields the equilibrium path beliefs: $\beta(\hat{s}, 0) = \beta(s_g^*(\underline{p}), 0) = 1$ and $\beta(s_b^*(\underline{p}), T_b(\underline{p})) = \beta(s_b^*(\bar{p}), T_b(\bar{p})) = 0$. In addition, with the monotonic beliefs property, we know that $\beta(s_1, T_1) = 0$ if $T_1 > \min\{T_b(\underline{p}), T_b(\bar{p})\}$. The beliefs for all other first-period records lie off the equilibrium path and as such are not tied down. Because the good incumbent does not use positive direct transfers in equilibrium, we impose the following off-equilibrium path beliefs: $\beta(s_1, T_1) = 0$ for all $T_1 > 0$.

Now notice that \hat{s} is the greatest subsidy level that the good incumbent could choose such that the bad incumbent would not have a strict incentive to mimic her first-period subsidy and direct transfer decision in either state. For all $s < \hat{s}$, the bad incumbent prefers the first-period record that maximizes her one period utility $(s_b^*(p), T_b(p))$ to the first-period record $(s, 0)$ in both states. For all $s > \hat{s}$, the bad incumbent prefers $(s, 0)$ to $(s_b^*(p), T_b(p))$ in at least one state. Therefore, we impose the following additional off-equilibrium path beliefs: $\beta(s, 0) = 0$ for all $s > \hat{s}$ and $\beta(s, 0) = 1$ for all $s \leq \hat{s}$.

Recall, the domestic citizen's optimal strategy is to reelect the incumbent with probability $\lambda(\beta(s_1, T_1))$ where $\lambda(\cdot)$ is the cumulative distribution function from which challenge's initial reputation is drawn. Given these beliefs, we know how the domestic citizen will behave. He will reelect the incumbent with probability one if he observes any first-period record $(s, 0)$ with $s \leq \hat{s}$. He will reelect the incumbent with probability zero if he observes any first-period record $(s, 0)$ with $s > \hat{s}$ or (s, T_1) with $T_1 > 0$.

The above strategy profile constitutes an equilibrium if no player has a strict incentive to deviate. Given the beliefs, we have shown that the domestic citizen has no incentive to deviate. Hence, we only need to examine the behavior of the incumbent.

We begin by considering the bad incumbent. If she were to choose a first-period record $(s, 0)$ with $s \leq \hat{s}$, she would be reelected with probability one in either state. However, by the definition of \hat{s} and because $u_b(s, 0, p) + \delta u_b^*(p)$ is strictly increasing in s regardless of the inverse demand function's realization, she would be strictly worse off than if she were to have chosen $(s_b^*(p), T_b(p))$. If she were to choose a first-period record $(s, 0)$ with $s > \hat{s}$, she would not be reelected and would only receive her one period utility. Hence, she could not be any better off from choosing $(s, 0)$ with $s > \hat{s}$ over $(s_b^*(p), T_b(p))$ in either state, because $(s_b^*(p), T_b(p))$ yields her maximal one period utility. Therefore, the bad incumbent has no incentive to deviate in either state.

We now consider the good incumbent. Suppose the inverse demand function is $p(Q) = \underline{p}(Q)$. If she were to deviate from $s_g^*(\underline{p}) = 0$, she could do so by choosing a subsidy level $s \in (0, \hat{s}]$. This would leave the probability that she is reelected unchanged, but would lead to a lower payoff because $u_g(s, 0, \underline{p}) + \delta u_g^*(\underline{p})$ is decreasing in s for all $s \geq 0$. If she were to deviate from $s_g^*(\underline{p}) = 0$, she could also do so by choosing a subsidy level $s \in (\hat{s}, \infty)$. This would make the probability that she is reelected equal to zero. Given that she would only receive her first-period utility, the optimal subsidy would be $s_g^*(\underline{p})$ which yields her maximal one period utility when the inverse demand function is $p(Q) = \underline{p}(Q)$. Therefore, the good incumbent has no incentive to deviate in state \underline{p} .

Now suppose the inverse demand function is $p(Q) = \bar{p}(Q)$ and that \hat{s} satisfies condition (T1). If the good incumbent were to deviate from \hat{s} , she could do so by choosing a subsidy level $s \in [0, \hat{s})$. This would leave the probability that she is reelected unchanged, but would lead to a lower payoff because $u_g(s, 0, \bar{p}) + \delta u_g^*(\bar{p})$ is increasing in s for all $s < s_g^*(\bar{p})$. If she were to deviate from \hat{s} , she could also do so by choosing a subsidy level $s \in (\hat{s}, \infty)$. This would make the probability she is reelected equal to zero. Given that she would only receive her first-period utility, the

optimal subsidy would be $s_g^*(\bar{p})$ which yields her maximal one period utility in state \bar{p} . However, from (T1) we know that she is better off by choosing $(\hat{s}, 0)$ and being reelected with probability one. Therefore, the good incumbent has no incentive to deviate when in state $\bar{p}(Q)$.

This completes the proof for the first part of Proposition 1. That is, we have shown that the above strategy profile constitutes a separating equilibrium. However, there are generally many other separating equilibria. This is the case whenever (T1) does not hold with strict equality at \hat{s} . In this situation, other separating equilibria take similar form as the one above except that the good incumbent in state \bar{p} chooses subsidy level $\hat{\hat{s}}$ where $\hat{\hat{s}} < \hat{s}$ and still satisfies (T1). These equilibria have different beliefs, namely that $\beta(s, 0) = 0$ for all $s > \hat{\hat{s}}$ and $\beta(s, 0) = 1$ for all $s \leq \hat{\hat{s}}$.

We now show that our \hat{s} separating equilibrium is the only strategy profile which satisfies the Intuitive Criterion. That is, any $\hat{\hat{s}}$ separating equilibria are eliminated by the Intuitive Criterion.

Consider a $\hat{\hat{s}}$ separating equilibrium: the good incumbent chooses $(\hat{\hat{s}}, 0)$ when $p(Q) = \bar{p}(Q)$ and $(s_g^*(\underline{p}), 0)$ when $p(Q) = \underline{p}(Q)$ and that bad incumbent chooses $(s_b^*(\underline{p}), T_b(\underline{p}))$ in both states. The beliefs are: $\beta(s, 0) = 0$ for all $s > \hat{\hat{s}}$, $\beta(s, 0) = 1$ for all $s \leq \hat{\hat{s}}$, and $\beta(s, T) = 0$ for all $T > 0$. Suppose the domestic citizen observes the first-period record $(\tilde{s}, 0)$ where \tilde{s} is such that $\hat{\hat{s}} < \tilde{s} < \hat{s}$. According to the beliefs, the domestic must assign probability zero to the event that this first-period record is sent by a good incumbent. However, by the definition of \hat{s} is it impossible that a bad incumbent chooses this first-period record. The Intuitive Criterion states that we must then assign probability one to the event that this first-period record is sent by a good incumbent. With this revised belief, the good incumbent in state \bar{p} has a strict incentive to deviate from $\hat{\hat{s}}$ to \tilde{s} . Hence, the $\hat{\hat{s}}$ separating equilibrium fails to satisfy the Intuitive Criterion. ■

3.3.2 Partial-Pooling Equilibria

Our second proposition captures the essence of the political economy critique of strategic export promotion: if the incumbent's initial reputation is sufficiently high, then the bad incumbent will use the subsidy to benefit the special interest and refrain from using the direct transfer even though the direct transfer yields the special interest a greater one period benefit.

Proposition 2 (Abuse of Strategic Export Subsidies) *There exists a $\hat{\theta}_I$ such that for all $\theta_I > \hat{\theta}_I$, there exists a partial-pooling equilibrium in which the good incumbent subsidizes the domestic firm at $s_g^*(\underline{p})$ in state \underline{p} and $s_g^*(\bar{p})$ in state \bar{p} , while the bad incumbent refrains from using the direct transfer, but always subsidizes the domestic firm at $s_g^*(\bar{p})$.*

Proof. If neither incumbent implements a positive first-period direct transfer and the good incumbent always chooses the subsidy level which maximizes her one period utility given the state, while the bad incumbent always chooses $s_g^*(\bar{p})$, then the domestic citizen's equilibrium path beliefs are:

$$\begin{aligned}\beta(s_g^*(\bar{p}), 0) &= \frac{\theta_I \phi^s}{\theta_I \phi^s + (1 - \theta_I)}, \\ \beta(s_g^*(\underline{p}), 0) &= 1.\end{aligned}$$

We now define the function $h : [0, 1] \rightarrow [0, 1]$ as

$$h(\theta) = \frac{\theta \phi^s}{\theta \phi^s + (1 - \theta)}.$$

Notice that h is continuous and increasing in θ and that $h(1) = 1$. Let $\hat{\theta}_I$ be defined as the value of θ such that

$$\lambda(h(\theta)) = \max \left\{ 1 - \frac{u_g^*(\bar{p}) - u_g(s_g^*(\underline{p}), 0, \bar{p})}{\delta u_g^*(\bar{p})}, \frac{u_b^*(\bar{p}) - u_b(s_g^*(\bar{p}), 0, \bar{p})}{\delta u_b^*(\bar{p})}, \frac{u_b^*(\underline{p}) - u_b(s_g^*(\bar{p}), 0, \underline{p})}{\delta u_b^*(\underline{p})}, 1 - \frac{u_b(s_g^*(\bar{p}), 0, \bar{p}) - u_b(s_g^*(\underline{p}), 0, \bar{p})}{\delta u_b^*(\bar{p})}, 1 - \frac{u_b^*(s_g(\bar{p}), 0, \underline{p}) - u_b(s_g^*(\underline{p}), 0, \underline{p})}{\delta u_b^*(\bar{p})} \right\}.$$

By assumptions (A1) and (A2) and the definitions of $u_g(\bar{p})$, $u_b(\bar{p})$, and $u_b(\underline{p})$, the value $\hat{\theta}_I$ exists and is in the interval $(0, 1)$.

We now show that, for $\theta_I > \hat{\theta}_I$, the good incumbent's always behaving to maximize her first-period utility and a bad incumbent's always choosing $s_g^*(\bar{p})$ is an PBNE with off-equilibrium path beliefs $\beta(s, 0) = 0$ for all $s \neq s_g^*(\underline{p}), s_g^*(\bar{p})$ and $\beta(s, T) = 0$ for all $T > 0$. We first check to see that the good incumbent has no incentive to deviate. When $p(Q) = \underline{p}(Q)$, it is clear that she will choose $(s_g^*(\underline{p}), 0)$ because this maximizes her first-period utility and gets her reelected with probability one. When $p(Q) = \bar{p}(Q)$, the payoff from $(s_g^*(\bar{p}), 0)$,

$$u_g(s_g^*(\bar{p}), 0, \bar{p}) + \delta \lambda(h(\theta_I)) u_g^*(\bar{p}).$$

The definition of $\hat{\theta}_I$ guarantees that this exceeds the payoff from choosing $(s_g^*(\underline{p}), 0)$ which is $u_g(s_g^*(\underline{p}), 0, \bar{p}) + \delta u_g^*(\bar{p})$. Furthermore, the payoff from $(s_g^*(\bar{p}), 0)$ exceeds the payoff from $(s, 0)$ for all $s \neq s_g^*(\underline{p}), s_g^*(\bar{p})$ which is $u_g(s, 0, \bar{p})$. To see this, notice that $(s_g^*(\bar{p}), 0)$ uniquely maximizes the good incumbent's one period utility when $p(Q) = \bar{p}(Q)$. Therefore, there is no strict incentive for the good incumbent to deviate.

Next, we check to see that the bad incumbent has no incentive to deviate from $(s_g^*(\underline{p}), 0)$ in both states of the world. When $p(Q) = \underline{p}(Q)$, her payoff from choosing

$(s_g^*(\bar{p}), 0)$ is

$$u_b(s_g^*(\bar{p}), 0, \underline{p}) + \delta \lambda(h(\theta_I)) u_b^*(\underline{p}).$$

The definition of $\hat{\theta}_I$ guarantees that this exceeds the payoff from choosing $(s_g^*(\underline{p}), 0)$ which is $u_b(s_g^*(\underline{p}), 0, \underline{p}) + \delta u_b^*(\underline{p})$. Furthermore, the payoff from $(s_g^*(\bar{p}), 0)$ exceeds the payoff from $(s, 0)$ for all $s \neq s_g^*(\underline{p}), s_g^*(\bar{p})$ which $u_b(s, 0, \underline{p})$. To see this notice that $(s_b^*(\underline{p}), T_b(\underline{p}))$ uniquely maximizes the bad incumbent's one period utility which, by the definition of $\hat{\theta}_I$, is strictly less than her payoff from $(s_g^*(\bar{p}), 0)$. A similar argument is true when $p(Q) = \bar{p}(Q)$. Therefore, there is no strict incentive for the bad politician to deviate. ■

3.3.3 Discussion

Proposition 1 tells us that when the domestic citizen is uncertain about the motivations of the incumbent and is uncertain about whether a positive subsidy to the domestic firm is in his interest, the good incumbent of state \bar{p} will attempt to separate herself from the bad incumbent of state \underline{p} and \bar{p} by decreasing her subsidy to the domestic firm. When the good incumbent of state \bar{p} decreases her subsidy away from $s_g^*(\bar{p})$, she decreases each bad incumbent's payoff from mimicking her and makes it more attractive for both bad incumbents to fully separate by using the direct transfer. This is because regardless of the state, the bad incumbent's per period utility function is strictly increasing in the subsidy level.

The highest subsidy level which makes both bad incumbents weakly indifferent between mimicking the good incumbent of state \bar{p} and fully separating is \hat{s} . Therefore, if the good incumbent chooses any the subsidy level weakly less than \hat{s} , she will be able to separate herself from the bad incumbents. The subsidy level \hat{s} is the least costly

way for the good incumbent of state \bar{p} to separate herself from the bad incumbents because the her payoff in state \bar{p} is strictly increasing in the subsidy level for all subsidies below $s_g^*(\bar{p})$. Moreover, this feature of the good incumbent's payoff in state \bar{p} with respect to the subsidy level means that if the good incumbent does not find it beneficial to separate by choosing \hat{s} then no separating equilibrium exists.

When the good incumbent of state \bar{p} is unable to separate herself from the bad incumbents, the possibility that subsidies are used to disguise income transfers to the special interest arises. In particular, Proposition 2 tells us that the bad politician in state \underline{p} will choose a positive subsidy as a means of transferring income to the special interest at the expense of the domestic citizen. This occurs when the incumbent has a sufficiently high initial reputation, i.e., when the domestic citizen believes with high probability that the incumbent is good. To understand why, notice that when the incumbent has a high initial reputation, both good and bad incumbents stand a strong chance of being reelected and therefore have a high expected second-period utility as long as they do not appear to as the bad incumbent to the domestic citizen at the time of the election. When the initial reputation is sufficiently high the bad incumbent's expected second-period gain utility associated with mimicking the good incumbent outweighs her utility loss from forgoing a first-period record that yields a greater first period payoff. That is, the reputational cost of choosing a first-period record other than that chosen by the good incumbent in state \bar{p} is sufficiently high that it is not in the bad incumbent of state \underline{p} 's interest to do so.

The existence of this partial-pooling equilibrium is the “cause for grave concern” which motivates Grossman's warning regarding the dangers of any scheme of targeted export promotion. That is, some politicians will implement an export subsidy when such a policy benefits the domestic citizen, while other politicians will implement an export subsidy when such a policy harms the domestic citizen. However,

as there is no way for the domestic citizen to distinguish the politicians, nothing can prevent bad politicians from abusing export subsidies. Therefore, from the domestic citizen's point of view schemes of target export promotion should be avoided.

Crucial to this argument is that the domestic citizen has both policy uncertainty and politician uncertainty. First, without policy uncertainty, the domestic citizen would know when subsidies are warranted. With this knowledge, the domestic citizen could always punish incumbents who inappropriately implement a subsidy by voting them out of office. Hence, bad politicians would not be able to use subsidies as means transferring income to the special interest while maintaining a positive reelection probability. Second, without politician uncertainty, the domestic citizen would know which incumbents act in their interest. Hence, bad incumbents would be voted out of office regardless of their first-period records.

3.4 A Continuum of Cournot Industries

We now extend our model to the case where there is a continuum of industries $[0, 1]$, each identical to the single industry in the previous section. However, we add the following wrinkle to the subsidy information available the domestic citizen at the time of voting: he can only observe that at least one firm is being subsidized or that no firm is being subsidized.

Suppose that the fraction of industries in which subsidies benefit the domestic citizen, k , is drawn from some cumulative distribution function $F(k)$. That is, k is the fraction of industries in which the inverse demand function is $\bar{p}(Q)$. We assume that the cumulative distribution function $F(k)$ is smooth and increasing with the property that $F(0) = 0$.

Consider the following strategy for the incumbent. If $k \geq k^*$, the good in-

cumbent implements subsidy $s_g^*(\bar{p})$ in industries in with $p(Q) = \bar{p}(Q)$ implements a zero subsidy (i.e., subsidy $s_g^*(\underline{p}) = 0$) in industries with $p(Q) = \underline{p}(Q)$. Otherwise, the good incumbent implements a zero subsidy in all industries. The bad incumbent implements subsidy $s_g^*(\bar{p})$ in all industries regardless of the inverse demand function's realization.

Our third proposition concerns the existence of partial-pooling equilibria when there is a continuum of identical industries.

Proposition 3 *There exists a $\tilde{\theta}_I$ s.t. for all $\theta_I > \tilde{\theta}_I$, there exists at least one partial-pooling equilibrium of the above form characterized by $k^*(\theta)$ if the following condition is satisfied:*

$$u_g(\bar{p}) > u_g(0, 0, \bar{p}) + \delta u_g(\bar{p}) \quad (T3)$$

Proof. If the incumbent follows the above strategy, then the domestic citizen's equilibrium path beliefs take simple form. If he observes no subsidy and a zero direct transfer, he has belief $\beta(N, 0, k^*) = 1$ and reelects the incumbent with probability one. If he observes someone being subsidized and a zero direct transfer, he has the belief $\beta(Y, 0, k^*)$ given by

$$\beta(Y, 0, k^*) = \frac{\theta_I \Pr(k \geq k^*)}{\theta_I \Pr(k \geq k^*) + (1 - \theta_I)} = \frac{\theta_I (1 - F(k^*))}{\theta_I (1 - F(k^*)) + (1 - \theta_I)},$$

and reelects the incumbent with probability $\lambda(\beta(Y, 0, k^*))$. From the monotonic beliefs property, we know that $\beta(Y, T, k^*) = \beta(N, T, k^*) = 0$ for all $T > 0$. Hence if any direct transfers are observed, the domestic citizen will not reelect the incumbent.

If the good incumbent were to subsidize a single industry, she would generate the belief $\beta(Y, 0, k^*)$ regardless of the subsidy level she chooses. Therefore, she would choose the subsidy level $s_g^*(\bar{p})$ which maximizes her one period utility from that in-

dustry without changing the probability with which she is reelected. Hence, the good incumbent will behave according to the above strategy if the following two conditions hold:

$$(1 + \delta \lambda(\beta(Y, 0, k^*))) [k u_g^*(\bar{p}) + (1 - k) u_g^*(\underline{p})] - k u_g(0, 0, \bar{p}) - (1 - k) u_g^*(\underline{p}) - \delta [k u_g^*(\bar{p}) + (1 - k) u_g^*(\underline{p})] \geq 0 \quad \forall k \geq k^* \quad (12)$$

$$(1 + \delta \lambda(\beta(Y, 0, k^*))) [k u_g^*(\bar{p}) + (1 - k) u_g^*(\underline{p})] - k u_g(0, 0, \bar{p}) - (1 - k) u_g^*(\underline{p}) - \delta [k u_g^*(\bar{p}) + (1 - k) u_g^*(\underline{p})] < 0 \quad \forall k < k^* \quad (13)$$

The first condition states that, whenever $k \geq k^*$, the good incumbent prefers to subsidize the industries which deserve subsidies in the first period and be reelected with probability $\lambda(\beta(Y, 0, k^*))$ than to not subsidize any industry in the first period and be reelected with probability one. The second condition states that, whenever $k < k^*$, the good incumbent prefers to not subsidize any industry and be reelected with probability one.

The right-hand sides of conditions (12) and (13) simplify to

$$\begin{aligned} & ((1 - \delta[1 - \lambda(\beta(Y, 0, k^*))]) [u_g^*(\bar{p}) - u_g^*(\underline{p})] - u_g(0, 0, \bar{p}) + u_g^*(\underline{p})) k \\ & - \delta[1 - \lambda(\beta(Y, 0, k^*))] u_g^*(\underline{p}). \end{aligned} \quad (14)$$

Notice that expression (14) is linear in k and is negative when $k = 0$. Furthermore, notice that expression (14) is strictly increasing in k if $(1 - \delta[1 - \lambda(\beta(Y, 0, k^*))]) u_g^*(\bar{p}) - u_g(0, 0, \bar{p}) > 0$. Condition (T3) guarantees this to be the case. Hence, there exists a $k > 0$, which we denote \hat{k} , such that expression (14) equals zero. We can express \hat{k}

as a function of k^* as follows:

$$\hat{k}(k^*) = \frac{\delta[1 - \lambda(\beta(Y, 0, k^*))]u_g^*(\underline{p})}{(1 - \delta[1 - \lambda(\beta(Y, 0, k^*))])[u_g^*(\bar{p}) - u_g^*(\underline{p})] - u_g(0, 0, \bar{p}) + u_g^*(\underline{p})} \quad (15)$$

The above strategy will be an equilibrium only if $\hat{k}(k^*) = k^*$, otherwise the good incumbent will have a strict incentive to deviate for some value of k . Notice the $\hat{k}(k^*)$ is continuous because belief $\beta(Y, 0, k^*)$ is continuous and because the cumulative distribution functions $F(k)$ and $\lambda(\cdot)$ are continuous. Therefore, we know $\hat{k}(k^*)$ has a fixed point in $(0,1)$ if $\hat{k}(0) > 0$ and if $\hat{k}(1) < 1$. Condition (T3) guarantees this to be the case. Therefore there exists a k^* such that $\hat{k}(k^*) = k^*$.

We now find the values of θ_I such that the bad incumbent has no incentive to deviate. If the bad incumbent were to subsidize a single industry without using a direct transfer, she would generate the belief $\beta(Y, 0, k^*)$ regardless of the subsidy level she chooses. Therefore, if she were to subsidize a single industry, she would be better off subsidizing all industries as her single industry per period utility is increasing in the subsidy level in both states of the world. Moreover, she would choose the subsidy level $\bar{s}(p)$ which yields the domestic citizen a zero one period utility from a single industry in state p when the direct transfer equals zero. Formally, we define $\bar{s}(p)$ as

$$\bar{s}(p) = s \quad s.t. \quad \alpha \pi(s, p) - s q(s, p) = 0.$$

Hence, the bad incumbent will behave according to the above strategy if the following condition holds:

$$\begin{aligned} k u_b(\bar{s}(\bar{p}), 0, \bar{p}) + (1 - k) u_b(\bar{s}(\underline{p}), 0, \underline{p}) + \delta \lambda(\beta(Y, 0, k^*)) [k u_b^*(\bar{p}) + (1 - k) u_b^*(\underline{p})] \\ \geq k u_b^*(\bar{p}) + (1 - k) u_b^*(\underline{p}) \quad \forall k. \end{aligned} \quad (16)$$

The above condition states that the bad incumbent prefers to forgo first-period direct transfers and subsidize all industries at the subsidy levels which leaves the domestic citizen zero first-period utility and gets her reelected with probability $\lambda(\beta(Y, 0, k^*))$ than to not subsidize any industry and be reelected with probability one.

Condition (16) can be rewritten as

$$\lambda(\beta(Y, 0, k^*)) = \frac{k u_b^*(\bar{p}) + (1 - k)u_b^*(\underline{p}) - k u_b(\bar{p}, 0, \bar{p}) + (1 - k)u_b(\bar{s}(\underline{p}), 0, \underline{p})}{\delta [k u_b^*(\bar{p}) + (1 - k) u_b^*(\underline{p})]} \forall k. \quad (17)$$

Because $\lambda(\theta)$ is a continuous and strictly increasing cumulative distribution function, it has an inverse. Hence, we can further manipulate equation (17) as follows:

$$\beta(Y, 0, k^*) \geq \lambda^{-1} \left(\max_k \frac{k u_b^*(\bar{p}) + (1 - k)u_b^*(\underline{p}) - k u_b(\bar{p}, 0, \bar{p}) + (1 - k)u_b(\bar{s}(\underline{p}), 0, \underline{p})}{\delta [k u_b^*(\bar{p}) + (1 - k) u_b^*(\underline{p})]} \right)$$

$$\frac{\theta_I (1 - F(k^*))}{\theta_I (1 - F(k^*)) + (1 - \theta_I)} \geq \lambda^{-1}(\cdot)$$

$$\theta_I \geq \frac{\lambda^{-1}(\cdot)}{(1 - F(k^*)) + F(k^*) \lambda^{-1}(\cdot)} \quad (18)$$

Equation (15) and (18) together define $\tilde{\theta}_I$. ■

3.5 Conclusion

We have presented a model which formalizes the political economy critique against strategic export promotion. When governments engage in such trade policy, it is possible that incumbents will abuse subsidies and use them as a disguised transfer mechanism to a special interest. Even if subsidies are not used in a corrupt manner, it is possible that incumbents will use subsidies in a suboptimal manner to build a

reputation as a good politician. These problems are further magnified if there are many industries which are eligible for export subsidies.

Appendix

Proof. [Proof: Lemma 2]

Disregard the non-negative constraint on the subsidy level. From equation (6), we know that the optimal subsidy from the domestic citizen's point of view is zero when the following holds:

$$\alpha \left[p' q^* \frac{dq_f^*}{ds} \right] + (\alpha - 1) q^* = 0 \quad (19)$$

We can find the dq_f^*/ds as a function of the primitives as follows. Obtain the following first order conditions from the domestic and foreign firm profit maximization problems:

$$\frac{d\pi}{dq} = p' q^* + p - c + s = 0 \quad \text{and} \quad (1')$$

$$\frac{d\pi_f}{dq_f} = p' q_f^* + p - c_f = 0 . \quad (2')$$

Using the implicit function theorem and equations (1') and (2'), we obtain

$$\frac{dq_f^*}{ds} = \frac{1 + \nu(q_f^*/Q^*)}{p'(3 + \nu)} , \quad (20)$$

where $\nu = (p''/p')Q^*$, i.e., the relative convexity the inverse demand function at the Cournot equilibrium.

Substituting (20) into (19) and isolating α , yields the following condition which

characterizes the domestic citizens' benefit from a positive subsidy:

$$\alpha^* = \frac{3 + \nu}{4 + \nu(1 + (q_f^*/Q^*))}. \quad (21)$$

When $\alpha > \alpha^*$, a positive subsidy benefits domestic citizens. When $\alpha \leq \alpha^*$, a non-positive subsidy benefits domestic citizens. In addition to (21), the strategic substitutes assumption provides another condition relating ν to foreign firm market share (q_f^*/Q^*) ,

$$\nu > \underline{\nu} \equiv \max \left\{ -\frac{1}{(q_f^*/Q^*)}, -\frac{1}{1 - (q_f^*/Q^*)} \right\}. \quad (22)$$

Conditions (20) and (21) along with the fact that $\alpha \in (0, 1)$, provide a relationship between the domestic citizens' share of profits, the relative convexity of the inverse demand function, and the foreign firm's market share to the domestic citizen's preference for a positive subsidy.

We define, $\underline{\alpha}$ and $\bar{\alpha}$ in the following manner. Holding fixed (q_f^*/Q^*) ,

$$\underline{\alpha} \equiv \min \left\{ \lim_{\nu \rightarrow \underline{\nu}} \alpha^*, \lim_{\nu \rightarrow \infty} \alpha^* \right\} \quad (23)$$

$$\bar{\alpha} \equiv \max \left\{ \lim_{\nu \rightarrow \underline{\nu}} \alpha^*, \lim_{\nu \rightarrow \infty} \alpha^* \right\} \quad (24)$$

Taking the limits of $\underline{\alpha}$ and $\bar{\alpha}$ as the foreign firm's market share goes to zero and unity, establishes that $\underline{\alpha} \in (\frac{1}{2}, \frac{3}{4}]$ and $\bar{\alpha} \in (\frac{3}{4}, 1]$.

We now turn to the relationship between marginal costs of the domestic and foreign firms and the foreign firm's market share. To see how Subtracting (2') from (1') and some algebraic manipulation yields the intuitive result that $q^* \geq q_f^*$ as $c_f \geq c$.

Hence, we have that

$$c_f > c \Rightarrow (q^*/Q^*) < \frac{1}{2} \quad \text{and} \quad c_f < c \Rightarrow (q^*/Q^*) > \frac{1}{2}. \quad (25)$$

This provides some insight into how the α^* line shifts as costs change. ■

References

- Brander, James. "Strategic Trade Policy," in: G. Grossman and K. Rogoff, eds., *Handbook of International Economics*, vol. 3. Amsterdam: Elsevier Science, 1995.
- Brander, James A. and Barbara Spencer. "Export Subsidies and International Market Share Rivalry." *Journal of International Economics* 18 (1985): 83-100.
- Coate, Stephen and Stephen Morris. "On the Form of Transfers to Special Interests." *Journal of Political Economy* 103 (1995): 1210-1235.
- Eaton, Jonathan and Gene Grossman. "Optimal Trade and Industrial Policy Under Oligopoly." *Quarterly Journal of Economics* 101 (1986): 383-406
- Grossman, Gene., "Strategic Export Promotion: A Critique," in: P. Krugman, ed., *Strategic Trade Policy and the New International Economics*. Cambridge, MA: MIT Press, 1986.
- Neary, J. Peter. "Cost Asymmetries in International Subsidy Games: Should Governments Help Winners or Losers?" *Journal of International Economics* 37 (1994): 197-218.