



KELLEY SCHOOL OF BUSINESS

INDIANA UNIVERSITY

Department of Business Economics and Public Policy

Astroturf Lobbying

Thomas P. Lyon

John W. Maxwell*

Kelley School of Business

Indiana University

Bloomington, IN 47405

July 2002

Abstract

We study three corporate non-market strategies designed to influence the lobbying behavior of other special interest groups: 1) “astroturf,” in which the firm covertly subsidizes a group with similar views to lobby when it normally would not, 2) the “bearhug,” in which the firm overtly subsidizes the lobbying activities of another interest group, and 3) self-regulation, in which the firm voluntarily limits the potential social harm from its activities. All three strategies can be used to reduce the informativeness of lobbying, and all reduce the welfare of the public decision maker. We show that the decision maker would benefit by requiring the public disclosure of funds spent on astroturf lobbying, but the availability of alternative influence strategies limits the impact of such a policy.

JEL Numbers: L51, Q28

Keywords: Special Interest Group, Lobbying, Corporate Strategy

*We would like to thank participants in seminars at Stanford University, Cambridge University, Katholieke University of Leuven, the University of Bonn, the 21st Rutgers Workshop on Public Utility Regulation, and the Second World Congress of the Association of Environmental and Resource Economists for their helpful comments.

1 Introduction

The role of interest groups in politics has held a long-standing fascination for political economists. In the 1780s, James Madison famously warned of the power of “factions” in *The Federalist*, while nearly two hundred years later Mancur Olson and George Stigler elevated the study of interest group politics to an important subfield within economics.¹ Pioneering theoretical work in the Chicago School tradition treated interest group “pressure” as a production function, smooth and twice continuously differentiable.² In this framework, interest groups compete to apply more pressure in a game where rival pressure inputs are strategic complements. More recently, theorists have been opening up the black box of political pressure to focus more explicitly on specific strategies such as campaign contributions or lobbying.³

Several recent papers shed new light on the role of lobbying in conveying “soft,” i.e. unverifiable, information to public decisionmakers.⁴ In these models, interest groups may be able to credibly transmit soft information if their preferences do not diverge too greatly from those of the decisionmaker. This recent work, however, typically does not distinguish firms from other special interest groups. We argue that in many lobbying situations, firms do indeed have preferences distinct from those of other groups. In particular, they often bear the costs of government policy but do not collect the benefits. This is especially true for policies dealing with externalities or the provision of public goods. In such circumstances, firms cannot credibly convey unverifiable information because their powerful bias towards weak policies is common knowledge among decisionmakers. Thus, existing models really cannot capture the role of the firm in lobbying games. Nevertheless, we show that firms can play an important role by influencing the lobbying behavior of other interest groups. The corporate strategies that accomplish this goal are the subject of our paper.

Most prominent among these strategies is the funding of “astroturf lobbying,” a term coined by Lloyd Bentsen, long-time Senator from Texas, to describe the artificial grassroots campaigns that are created by public relations (PR) firms.⁵ One such firm is Davies Communications, whose advertising says “Traditional lobbying is no longer enough. Today numbers count. To win in the hearing room, you must reach out to create grassroots support.

¹Olson (1965) elaborates a rational choice model of interest group action, while Stigler (1971) applies this approach to the study of regulation specifically.

²Key early papers include those of Peltzman (1976) and Becker (1983).

³Grossman and Helpman (2001) provide an excellent introduction to the more recent theoretical literature on interest group politics.

⁴See, for example, Lohmann (1993) and Krishna and Morgan (2001). These models, which build on the seminal work of Crawford and Sobel (1982), must be distinguished from models of the provision of “hard,” verifiable, information, as analyzed in papers such as Milgrom and Roberts (1986).

⁵Stauber and Rampton (1995), p. 79.

To outnumber your opponents, call the leading grassroots public affairs communications specialists.”⁶ Davies explains how his firm generates a “grassroots” letter-writing campaign through the use of telephone banks:

“We get them on the phone, and while we’re on the phone we say ‘Will you write a letter?’ ‘Sure.’ ‘Do you have time to write it?’ ‘Not really.’ ‘Could *we* write the letter for you? I could put you on the phone right now with someone who could help you write a letter. Just hold, we have a writer standing by’...If they’re close by we hand-deliver it. We hand-write it out on ‘little kitty cat stationery’ if it’s a little old lady. If it’s a business we take it over to be photocopied on someone’s letterhead. [We] use different stamps, different envelopes. Getting a pile of personalized letters that have a different look to them is what you want to strive for.”⁷

One example of astroturf lobbying is the group People for the West!, which characterizes itself as “a grassroots campaign supporting western communities.” In 1992, 96% of the group’s funding came from corporate sponsors such as NERCO Minerals, Cyprus Minerals, Chevron, and Hecla Mining, who have strong interests in maintaining the General Mining Act of 1872 that allows them to acquire and mine public lands at a cost of \$5 per acre. The chairman of PFW!, Bob Quick, is the national director of state legislative affairs for Asarco, a mining company.⁸ Another example is the Consumer Alliance, a Michigan-based nonprofit that opposes laws to lower the price of prescription drugs to Medicaid participants and other low-income citizens. A public relations firm called Bonner & Associates—funded by the Pharmaceutical Research and Manufacturers of America (PhRMA)—uses Consumer Alliance letterhead to solicit signatures in support of its positions.⁹

Astroturf lobbying relies on the covert nature of corporate sponsorship in achieving its effectiveness. On December 19, 1995, President Clinton signed into law the Lobbying Disclosure Act of 1995, establishing new registration and reporting requirements for lobbyists working for corporations, charities and other nonprofit organizations engaged in efforts to influence legislative and executive branch decisions. The 1995 Act was the first major legislation on lobbying in nearly 50 years, and was designed to provide transparency in the lobbying process. Early drafts of the Lobbying Disclosure Act included provisions requiring the registration of firms engaged in astroturf lobbying, and the reporting of the expenditures made on those actions. Those provisions, however, failed to make it out of committee. As

⁶Stauber and Rampton (1995), p. 90.

⁷Stauber and Rampton (1995), pp. 89-91.

⁸For further details, see Sanchez (1996).

⁹For more details, see Craig (2002).

the bill’s sponsor, Senator Carl Levin, testified before a House committee considering the bill:

“Every reference to grass roots lobbying — and even to paid efforts to stimulate artificial grass roots lobbying — has been deleted from the bill...I am personally disappointed that we were unable to do anything to address the issue of a form of grass-roots lobbying referred to as “astroturf” lobbying, in which lobbyists hire professional experts to run phone banks and generate mail in support of their efforts. In my view, these paid, professional astroturf campaigns bear nothing in common with the genuine grassroots activities...I ...hope that the House will reconsider the disclosure of such lobbying...”¹⁰

Thus, a significant and growing aspect of the lobbying process remains obscured from public view.

In the present paper we develop a formal model of the lobbying process, focusing on the role of special interest groups in transmitting information to decisionmakers. We use the model to study a variety of strategies corporations can use to influence the lobbying behavior of other interest groups, all of which share the feature that they reduce the flow of information to public decisionmakers. Our model provides clear support for public disclosure of corporate expenditures on astroturf lobbying efforts, as called for by Senator Levin. However, such disclosure is not a panacea. We also study two other corporate strategies that can impede the flow of information, even when their use is common knowledge among all participants in the lobbying process.

We coin the term “bear hug” to refer to a corporate strategy of embracing one’s opposition by *overtly* subsidizing its lobbying efforts. This undermines the opposition’s ability to transmit its information through costly signaling. For example, DeSimone and Popoff point out that

“It is also important to recognize that there can be a disparity of resources and information between business stakeholder groups that makes trust difficult to develop. This may sometimes require action to redress the balance. Since the Brent Spar incident—when opposition prevented Shell from disposing of a large oil storage platform at sea—the company has made space available for environmental groups to explain their point of view in educational and other materials that it has prepared.”¹¹

¹⁰Testimony of Senator Carl Levin, Committee on the Judiciary, Subcommittee on the Constitution, U.S. House of Representatives, September 7, 1995.

¹¹DeSimone and Popoff (2000), p. 165.

The third strategy we examine, self-regulation, is quite different from the other two strategies, in that it involves real changes in company operations that are designed to reduce the risks of social harm. If these actions are substantive enough, interest groups may decide that the further gains from lobbying are not enough to justify the costs, and they may eschew participation in the political process. Perhaps surprisingly, however, we show that self-regulation may also benefit the firm by inducing interest groups to participate more actively in the lobbying process.

The remainder of the paper is organized as follows. Section 2 presents a simple model of the lobbying process. Section 3 studies astroturf lobbying, while section 4 considers the bear hug. Section 5 addresses the effects of self regulation, and section 6 discusses extensions of our model to a setting with multiple interest groups. Section 7 concludes.

2 A Simple Model of Lobbying

Our basic model of lobbying is based on Grossman and Helpman (2001), and begins with two players, a government decisionmaker (DM) and a special interest group (SIG).¹² In this version of the model the firm has no active role in the lobbying process. We assume the existence of a proposal that affects the firm and requires the approval of the decisionmaker, who may impose a variety of requirements on its passage to ensure that it is socially beneficial. The proposal might be an application for planning approval of a new manufacturing facility, in which case the DM may require that the manufacturer install certain emissions control systems as a condition of operation. In a legislative context, the proposal might call for amendment of the General Mining Act of 1872, in which case the DM might require the use of auctions to allocate mining rights on public lands, ensuring that the fiscal impact of the Act is minimized. Alternatively, the proposal might be aimed at health care reform, in which case the DM might require state Medicaid programs to negotiate the lowest possible prices from pharmaceutical manufacturers. In each case, the DM's proposal gives the affected firms a powerful incentive to attempt to influence the policy process.

For ease of presentation we will focus on a decision marker's choice of stringency of a local planning permit to build a local manufacturing facility. It is important to note, however, that the modeling of the decisionmaker as a unitary actor does not limit its applicability to the planning context. Other authors, such as Lohmann (1993), have used unitary-actor models to represent a political leader who responds to the preferences of the median voter. We will discuss our model's implications for the legislative context as appropriate below.

The construction of a manufacturing plant may have social effects through a variety of

¹²See Grossman and Helpman (2001), chapter 5.

mechanisms, e.g. it may create jobs in the local community, it may affect the environmental quality of the surrounding community, and it may affect the health and safety of that community. These effects can be summarized by a variable $\theta \in \mathfrak{R}$, which represents the true state of the world. For simplicity, we will assume that the state of the world captures the net adverse social impact of the project, and can either be “low” or “high,” so $\theta \in \{\theta_L, \theta_H\}$.

The DM chooses a policy p that indicates the stringency of the regulatory response to the project. The DM is assumed to care about his or her constituency, perhaps because of reelection concerns. The DM’s preferences are represented by $G = -(p - \theta)^2$, which implies that the DM attempts to precisely match the policy to the state of the world. If the project is likely to have a highly adverse social impact on the local community, then the DM would favor setting a more stringent regulatory policy. Setting a policy that is higher than the true state is undesirable for the DM, because, for example, doing so might bring unnecessary economic hardship to the firm, which may in turn negatively affect employment in the local community. Setting too low a stringency is also undesirable for the DM, since community environmental, health, and safety conditions may be adversely affected. The DM’s prior belief is that either state of the world is equally likely. Without further information, the DM’s best policy decision is to

$$\max_p \frac{1}{2}[-(p - \theta_L)^2] + \frac{1}{2}[-(p - \theta_H)^2].$$

Consequently, under conditions of uncertainty the DM’s optimal decision is to set a moderately stringent policy of $p = (\theta_L + \theta_H)/2$ with $E(G) = -(\theta_H - \theta_L)^2/4$. We refer to a policy set at this level as the “average” policy.

The SIG is assumed to know the true state of the world.¹³ The SIG’s preferences are given by $U = -(p - \theta - \delta)^2 - l$, where δ represents the divergence between the SIG’s preferences and those of the DM, and l represents the cost to the SIG of lobbying the DM. Given this specification, the SIG always prefers a higher (lower) level of the policy p than does the DM when δ is positive (negative). We refer to δ as the SIG’s bias. The lobbying cost l may include not only the direct costs of lobbying, but also the cost of discovering the true state of the world, i.e., the true social impact of the firm’s project on the local community. The general form of the SIG’s utility function captures the assumption that the SIG cares about both the project’s social *and* economic effects on the local community. That is to say, even a positive biased SIG may prefer a less stringent policy to a more stringent policy if the true state of the world is low enough.

The location of manufacturing plants is often plagued by opposition from local residents

¹³This may reflect technical knowledge, e.g. regarding groundwater flow in regions of karst topography, or social knowledge, e.g. regarding local community preferences.

who proclaim that the plant can be built, but “not in my backyard.” While this may be a purely political phenomenon in some cases, in others it may reflect local knowledge of community preferences over the impacts of the project. In any event, it is natural to assume $\delta > 0$ in this situation, and we use this assumption in laying out the basic structure of the model. We begin our analysis with the case where $l = 0$, i.e., the (positive-biased) SIG knows the true state of the world and can costlessly lobby (report the state to) the DM. We examine the SIG’s incentives to report the true state of the world when the DM believes the SIG’s announcement. Since the SIG always prefers a higher level of policy than the DM, it naturally has no incentive to misreport when the state is $\theta = \theta_H$. Misreporting may be desirable however, if in $\theta = \theta_L$. In this case, the SIG misreports, i.e. reports θ_H , if its utility of obtaining θ_H in the low state exceeds its utility from reporting truthfully, that is, if:

$$-(\theta_L - \theta_L - \delta)^2 < -(\theta_H - \theta_L - \delta)^2$$

Thus, when $\theta = \theta_L$, the SIG misreports if

$$\delta > (\theta_H - \theta_L)/2. \tag{1}$$

Consider a case where condition (1) holds. This implies that the SIG has a large degree of bias, or alternatively, that the high and low states are relatively close together. In this case, the SIG will always report that $\theta = \theta_H$, regardless of the actual state of the world. Assuming the DM knows δ, θ_L , and θ_H , he will recognize the SIG’s incentives, and hence will not update his prior based on the SIG’s report. Thus, the DM sets $p = (\theta_L + \theta_H)/2$. If condition (1) fails to hold, then the SIG will report truthfully, and the DM will use the SIG’s report to set a policy of θ_L in the low state, and θ_H in the high state.¹⁴

Turn next to the case where lobbying is costly. Because the SIG is biased toward high levels of policy, it is particularly concerned about the possibility that the DM sets $p = \theta_L$ when the state is actually θ_H . Thus, the SIG is strongly motivated to incur the cost of lobbying when the state is θ_H , but may not find it worthwhile when $\theta = \theta_L$. Under certain conditions, which we explain below, there exists a Perfect Bayesian Equilibrium (henceforth, an “equilibrium”) in which the SIG only lobbies when $\theta = \theta_H$. In the equilibrium, the DM holds the belief that if the SIG lobbies then indeed $\theta = \theta_H$, and if the SIG fails to lobby then the state is θ_L . For this equilibrium to exist, the SIG must prefer to refrain from lobbying when $\theta = \theta_L$, i.e., $-(\theta_L - \theta_L - \delta)^2 \geq -(\theta_H - \theta_L - \delta)^2 - l$, or¹⁵

¹⁴Even when (1) fails, the truthful equilibrium is not unique. There always exists an equilibrium in which the DM distrusts the SIG’s information, and hence always sets the average policy. As a result, any signal by the SIG constitutes a best response. This equilibrium is not particularly interesting, however, and we will not consider it in the remainder of the paper.

¹⁵Note that with some rearranging of terms, the following expression reduces to (1) when $l = 0$.

$$l \geq \underline{l} \equiv (\theta_H - \theta_L)(2\delta - (\theta_H - \theta_L)). \quad (2)$$

At the same time, the SIG must be willing to incur the lobbying cost when the state is high, i.e. $-(\theta_H - \theta_H - \delta)^2 - l \geq -(\theta_L - \theta_H - \delta)^2$, which can be rewritten as

$$l \leq \bar{l} \equiv (\theta_H - \theta_L)(2\delta + \theta_H - \theta_L). \quad (3)$$

If both (2) and (3) hold, then the equilibrium described above exists; in the remainder of the paper, we will assume these conditions hold. Thus, a positive lobbying cost *aids* the SIG in truthful reporting by allowing it to express the intensity of its preferences. As we shall see in the subsequent sections, this result gives rise to a number of somewhat unexpected corporate strategies aimed at undermining the SIG's ability to express the intensity of its preferences.

Letting $a = (\theta_H - \theta_L)$, Figure 1 illustrates the values of l and a that give rise to truthful reporting by the SIG. The top line in the figure represents the combinations of l and a for which the SIG is just indifferent between lobbying when the true state of the world is θ_H and not lobbying in that state. Above this line, the SIG will choose not to incur the costs of lobbying even in the high state. The lower line traces out the combinations of l and a for which the SIG is just indifferent between lobbying in the low state (and falsely announcing θ_H) and not lobbying in the low state. For all combinations of l and a below the lower line, the SIG would strictly prefer to lobby in the low state (and announce θ_H). (Note that for $a \geq 2\delta$, the SIG would report truthfully in both states of the world, even if lobbying were costless.)

Consider the case of siting a new paper-making facility, which will release some volume of organochlorines into a river. The facility could use a number of alternative technologies for bleaching the pulp, which vary in their use of chlorine in the bleaching process and, thus, in the amount of organochlorines they release into the environment. A local environmental organization is concerned about organochlorine releases, since they result in the presence of trace amounts of dioxins—known carcinogens—in the river downstream from the plant. Suppose condition (1) holds and lobbying is costless. In this case, the environmental group will always participate in hearings about the plant, and it will argue that dioxins are highly toxic chemicals whose release should be avoided, regardless of the bleaching technology to be used and the quantity of releases involved. Since the group will always protest regardless of the firm's technology, its actions convey little about the intensity of its concerns about the technology. If it is costly for the group to participate in the hearings, however, then the net benefits of participation are small when dioxins are released in minute amounts, so the group

will eschew participation in that case. It will allocate its scarce lobbying resources to fighting the plant only when relatively large amounts of dioxins are likely to be released. Thus, when lobbying is costly and the local group does show up to participate in the proceedings, this is credible evidence that the harm from the plant’s dioxin releases is likely to be large, i.e. the true state is θ_H . In this case the decisionmaker can learn from the actions of the environmental group.¹⁶

An example in the legislative context would be a decision by lawmakers on whether to require the state to negotiate with pharmaceutical companies to obtain lower drug prices for Medicaid recipients. The consumer advocacy group Consumer Alliance might oppose such negotiations on the grounds that they might result in reduced choice in prescription drugs for senior citizens. If input from Consumer Alliance is solicited, it might oppose any proposal regardless of the extent to which it limited choice. If instead Consumer Alliance had to expend resources to mount a grass-roots campaign against draft legislation, then its decision to do so on any specific legislation might serve as a useful signal of the extent to which the legislation would limit choice.

3 Astroturf

We now introduce the firm as an active player in the game. Let the firm’s objective function be $F = -\beta p^2$, where $\beta > 0$. The parameter β can be interpreted as an efficiency parameter. Firms with large β ’s tend to be less efficient at adapting to more stringent policies. The structure of the firm’s objective function indicates that profits are strictly declining in the stringency of the DM’s policy, as is typical in economic models of regulation. This might be the case, for example, for the permitting requirements imposed on a proposed new manufacturing facility. The vast majority of the firm’s shareholders do not live in the local community, and hence are not directly affected by issues such as the availability of jobs within the community or environmental impacts of the plant. Assuming the DM is aware of the firm’s objectives, then it is easy to see that the firm is not a credible source of information regarding the state of the world: regardless of the true state, the firm has incentives to claim the state is θ_L .¹⁷

In this section, we consider the corporate strategy of “astroturf,” in which the firm subsidizes the lobbying cost of a special interest group after the firm learns the state of the world. These subsidy payments are made in states in which the special interest group would normally not lobby. This artificially-induced lobbying is called “astroturf lobbying.” This

¹⁶For further details on the issue of chlorine in the papermaking process, see Beckenstein *et al.* (1994).

¹⁷As a consequence, it is pointless for the firm to lobby the DM directly. It may, however, exert considerable influence over the actions of the other players.

strategy involves covertly supporting an interest group whose bias is negative; astroturf is thus a form of costly state falsification.¹⁸ As we noted in the Introduction, the most common examples of astroturfing involve the hiring of public relations or lobbying firms to create artificial grass-roots campaigns. In some cases these firms may subsidize the activities of legitimate grass-roots groups that are sympathetic to their clients' goals. The subsidies may be direct monetary payments, but they often involve providing free use of the firms' phone bank equipment and personnel. In the latter case, the employees of the public relations firm will pose as members of the grass-roots group when they make phone calls or send faxes.¹⁹

The *ex post* nature of the firm's subsidy payment is an important characteristic and distinguishes the strategy from the "bear hug" strategy, which we will examine in the following section. In many situations, the firm will know the true state of the world prior to making its project proposal. For example, the literature on environmental justice argues that firms take community characteristics and impacts into account when deciding where to site industrial plants.²⁰ In the context of health care reform, pharmaceutical companies presumably know in advance the true extent to which they will cut R&D spending if Medicaid reforms reduce the prices paid by the states for prescription drugs.

In our model it is the negative-biased group that is the natural ally of the firm, since the SIG's optimal policy outcome is more lenient than the DM's optimal policy in all states of the world.²¹ Such a group may place a greater value on the economic impacts of the firm's facility than does the decisionmaker. While the DM *does not* observe the firm's subsidy to the SIG, we will allow the DM to invest in auditing the SIG in order to determine whether or not the SIG has been subsidized.

We characterize conditions under which astroturf lobbying constitutes an equilibrium in our model. We begin by considering a single SIG with $U = -(p - \theta - \delta)^2 - l$, where $\delta < 0$. Suppose that lobbying is costless ($l = 0$), so the SIG can costless deliver a report on the true state of the world. Since the SIG always prefers a lower level of policy than does the DM, it has no incentive to misreport when the state is $\theta = \theta_L$. Misreporting may arise, however, if $\theta = \theta_H$, since the SIG may prefer to obtain a lower policy than θ_H even in the high state. The SIG will fail to report truthfully if

$$-(\theta_H - \theta_H - \delta)^2 < -(\theta_L - \theta_H - \delta)^2.$$

¹⁸Crocker and Tennyson (1997) study costly state falsification in the context of insurance, and show that the optimal insurance contract typically involves a strictly positive amount of falsification.

¹⁹For an example, see Craig (2002).

²⁰See Taylor (1992) or Greer and Harding (1993).

²¹Note that astroturf lobbying will not be used with a positive-biased SIG. Such a SIG chooses not to lobby in the low state, which leads the DM to set a low level of policy. Subsidizing the SIG to lobby in this state would induce the DM to set a stringent policy and make the firm worse off.

Thus, when $\theta = \theta_H$, the SIG will send a false report if

$$\delta \leq -(\theta_H - \theta_L)/2. \quad (4)$$

We can see from condition (4) that the SIG has an incentive to misreport if its bias is greater than half the distance between the two states. If condition (4) holds, then the SIG will always report θ_L and the DM's optimal response to the SIG's announcement will be to set the average policy, since the announcement is not credible. Paralleling our result in section 2, it is possible for the SIG to lobby credibly, even when condition (4) holds, if lobbying is costly. In this case, the SIG only lobbies when the state is low, since a policy mistake in this state is very costly to the SIG; if the state is high, however, the SIG may find it too costly to lobby. As a result, the DM can infer that the state is low when the SIG lobbies, and high when it does not lobby. For this equilibrium to exist, the SIG must prefer to refrain from lobbying when $\theta = \theta_H$, i.e., $-(\theta_H - \theta_H - \delta)^2 \geq -(\theta_L - \theta_H - \delta)^2 - l$, or

$$l \geq \underline{l} \equiv (\theta_H - \theta_L)(-2\delta - (\theta_H - \theta_L)). \quad (5)$$

Note that $\underline{l} > 0$ since $\delta < 0$.

At the same time, the SIG must be willing to incur the lobbying cost when the state is low, i.e. $-(\theta_L - \theta_L - \delta)^2 - l \geq -(\theta_H - \theta_L - \delta)^2$, which can be rewritten as

$$l \leq \bar{l} \equiv (\theta_H - \theta_L)(-2\delta + \theta_H - \theta_L). \quad (6)$$

If both (5) and (6) hold, then the equilibrium described above exists.

The question we wish to investigate is: Can the firm use astroturf lobbying to raise its expected payoff relative to its payoff when the SIG engages in truthful lobbying behavior? Recall that for astroturf lobbying to work the firm's subsidy to the negatively biased SIG must occur *ex post*, and be hidden from the DM. Although we assume that the DM cannot costlessly observe the firm's subsidy payment, it is clear from our discussion in the Introduction that policy makers are aware of the possibility of the astroturf lobbying strategy. Thus, we assume the DM can expend some resources in auditing the SIG's actions in an attempt to determine whether a subsidy did in fact occur. Let α denote the probability with which the DM conducts an audit, and τ denote the probability an audit, if conducted, generates conclusive information about whether a subsidy was conferred; with probability $1 - \tau$ the DM obtains no information when an audit is conducted. The cost of auditing is $c(\alpha)$, where $c'(\alpha) > 0$ and $c''(\alpha) > 0$. We will also assume that $\lim_{\alpha \rightarrow 0} c'(\alpha) = 0$ and $\lim_{\alpha \rightarrow 1} c'(\alpha) = \infty$, which assures an interior solution.

There are two possible types of equilibria with auditing, one in which astroturf does not occur and one in which it does.

The “No Astroturf” Equilibrium In this equilibrium, the DM believes (correctly) that if the SIG lobbies then the state is θ_L , and if the SIG does not lobby then the state is θ_H . To ensure these conditions hold, however, the DM must audit the SIG when it lobbies, in order to eliminate the firm’s incentives to astroturf. Assuming the firm does not engage in astroturf lobbying, the DM can infer correctly the state of the world, and sets the optimal policy for each state. Let the DM’s equilibrium audit probability in this case be α^{NA} . Thus, the DM’s expected payoff is $G^{NA} = -c(\alpha^{NA})/2$, since setting the correct policy generates an optimal utility of zero in both states.

Conditional on the DM’s audit policy, and the DM’s recognition that they are playing the “No Astroturf” equilibrium, the firm must prefer not to astroturf in state θ_H . (It need not engage in astroturf in state θ_L , as the SIG lobbies by assumption.) The firm’s profits if it does not astroturf are $\pi^{NA}(\theta_H) = -\beta\theta_H^2$. If it were to astroturf, its expected profits would be

$$\bar{\pi}^A(\theta_H) = \alpha^{NA}\tau(-\beta\theta_H^2) + (1 - \alpha^{NA}\tau)(-\beta\theta_L^2) - l.$$

Thus, with probability $\alpha^{NA}\tau$, the DM conducts an audit and the audit reveals that the firm engaged in astroturf; the DM then sets a high level of policy. With probability $1 - \alpha^{NA}\tau$ the DM obtains no new information, either because he does not audit or because the audit is uninformative; since the DM believes the “No Astroturf” equilibrium is being played, and has no evidence to the contrary, he sets a low level of policy.

A “No Astroturf” equilibrium requires $-\beta\theta_H^2 > \alpha^{NA}\tau(-\beta\theta_H^2) + (1 - \alpha^{NA}\tau)(-\beta\theta_L^2) - l$. This can be rewritten as

$$(1 - \alpha^{NA}\tau)\beta(\theta_H^2 - \theta_L^2) - l < 0. \tag{7}$$

In order to enforce the “No Astroturf” equilibrium, the DM must choose α^{NA} to make inequality (7) hold. This implies

$$\alpha^{NA} \geq \frac{\beta(\theta_H^2 - \theta_L^2) - l}{\tau\beta(\theta_H^2 - \theta_L^2)}. \tag{8}$$

Note that as τ becomes smaller, and the audit becomes less likely to be informative, the DM must audit with a higher probability. In fact, for small enough τ , the constraint that $\alpha^{NA} \leq 1$ becomes binding, and it becomes impossible for the DM to satisfy inequality (7). As a result, the firm finds it profitable to engage in astroturf, and the “No Astroturf” equilibrium does not exist.

The “Astroturf” Equilibrium Next we must consider the potential existence of an alternative equilibrium in which it is common knowledge that the DM does not audit enough to deter astroturf lobbying. In this equilibrium, the SIG always lobbies regardless of the state of the world: in the low state the SIG itself is motivated to lobby, while in the high state the firm pays the SIG to lobby. As a result, the DM always sets the average policy unless an audit catches the firm engaging in astroturf; in this case, the DM knows the state is θ_H and sets a stringent policy. We will assume that if the SIG does not lobby—which is an out-of-equilibrium event—then the DM believes the state must be θ_H , and sets $p = \theta_H$.²²

Let us consider the SIG’s optimal lobbying strategy in each state of the world. Suppose the state is $\theta = \theta_L$. The SIG obtains policy $p = (\theta_H + \theta_L)/2$ if it lobbies and policy $p = \theta_H$ if it does not. Lobbying is worthwhile if $-((\theta_H + \theta_L)/2 - \theta_L - \delta)^2 - l \geq -(\theta_H - \theta_L - \delta)^2$, which can be rewritten as

$$l \leq \bar{l}' \equiv (\theta_H - \theta_L)(3(\theta_H - \theta_L)/4 - \delta). \quad (9)$$

Now suppose the state is $\theta = \theta_H$. Again, the SIG obtains policy $p = (\theta_H + \theta_L)/2$ if it lobbies and policy $p = \theta_H$ if it does not. Lobbying is not worthwhile if

$$-(\theta_H - \theta_H - \delta)^2 > -((\theta_H + \theta_L)/2 - \theta_H - \delta)^2 - l.$$

This can be re-written as

$$l \geq \underline{l}' \equiv \frac{(\theta_H - \theta_L)}{2} \left[-2\delta - \frac{(\theta_H - \theta_L)}{2} \right]. \quad (10)$$

If both (9) and (10) hold, then in equilibrium the SIG’s optimal strategy is to lobby only in the low state (unless it is subsidized by the firm to lobby in the high state).²³ Figure 2 illustrates the curves \bar{l}' and \underline{l}' in relation to the curves \bar{l} and \underline{l} derived in section 2. The mathematical formulation of the curves differs now because in the simple model, if the SIG lobbies it expects the DM to set the policy it advocates, whereas in the astroturf equilibrium, lobbying yields only the average policy. As a result, lobbying is less productive for the SIG and the curves for the astroturf equilibrium are effectively “stretched” to the right, though they maintain the same general shape as the original curves. For the region where $l \in (\underline{l}', \bar{l}')$, the astroturf equilibrium exists if the firm finds it profitable to subsidize the SIG’s lobbying activity when the state is high.

²²This is consistent with the requirement of universal divinity, which requires placing all probability on the state in which the SIG would benefit most from deviating from equilibrium; in this case, the SIG benefits more from a deviation in the high state.

²³Note that $\underline{l}' > 0$ and $\bar{l}' > 0$ since $\delta < -(\theta_H - \theta_L)/2$.

We now examine whether it is profitable for the firm to engage in astroturf in the high state. If $\theta = \theta_H$ and the firm chooses to engage in astroturf, then

$$\bar{\pi}^A(\theta_H) = \alpha^A \tau(-\beta \theta_H^2) + (1 - \alpha^A \tau) \left[-\beta \left(\frac{\theta_H + \theta_L}{2} \right)^2 \right] - l.$$

The firm's expected profits reflect the fact that the stringent policy is imposed only if an audit reveals that astroturf lobbying occurred; this happens with probability $\alpha^A \tau$. Otherwise, the DM sets the average policy since he believes (correctly) that the "Astroturf" equilibrium is being played.

If the firm did not pay for astroturf lobbying when the state was $\theta = \theta_H$, then the SIG would not lobby. As noted above, this is out-of-equilibrium behavior, given that the DM believes they are playing the "Astroturf" equilibrium, and we assume that in this event the DM believes the state is θ_H , and sets $p = \theta_H$. As a result, the firm earns $\pi = -\beta \theta_H^2$. To ensure this deviation from equilibrium play does not occur, it must be the case that $\bar{\pi}^A(\theta_H) > -\beta \theta_H^2$. That is, an "Astroturf" equilibrium requires

$$\alpha^A \tau(-\beta \theta_H^2) + (1 - \alpha^A \tau) \left[-\beta \left(\frac{\theta_H + \theta_L}{2} \right)^2 \right] - l + \beta \theta_H^2 > 0, \quad (11)$$

which implies

$$\frac{1 - \alpha^A \tau}{4} \beta (3\theta_H + \theta_L)(\theta_H - \theta_L) > l. \quad (12)$$

This inequality must be consistent with conditions (9) and (10). The potentially binding constraint here is (9). Conditions (12) and (9) together require that

$$\frac{1 - \alpha^A \tau}{4} \beta (3\theta_H + \theta_L)(\theta_H - \theta_L) \geq (\theta_H - \theta_L)(3(\theta_H - \theta_L)/4 - \delta), \quad (13)$$

which can be rewritten as

$$\beta \geq \frac{4(3(\theta_H - \theta_L)/4 - \delta)}{(1 - \alpha^A \tau)(3\theta_H + \theta_L)}. \quad (14)$$

Thus, we obtain the following lemma.

Lemma 1 *For β satisfying inequality (14), the firm finds it profitable to fund the SIG to engage in astroturf lobbying.*

Lemma 1 states that if the firm's payoff function is sufficiently concave, then it is profitable to engage in the astroturf strategy, *i.e.*, to subsidize the SIG in the high state of the world even when it faces a positive probability of detection. In doing so, the firm benefits from obtaining the average policy in the high state (as long as an audit does not detect the

subsidy), although it does sacrifice the possibility of obtaining $p = \theta_L$ (obtaining the average policy instead) when the state of the world is low.

Finally, to determine whether an “Astroturf” equilibrium exists, we need to check whether the DM would prefer to deter astroturf and shift to the “No Astroturf” equilibrium. The DM’s expected utility in the “Astroturf” equilibrium is

$$\begin{aligned}
E(G^A) &= \frac{1}{2}[\alpha^A \tau(0) + (1 - \alpha^A \tau)(-\frac{\theta_H + \theta_L}{2} - \theta_L)^2] \\
&\quad + \frac{1}{2}[\alpha^A \tau(0) + (1 - \alpha^A \tau)(-\frac{\theta_H + \theta_L}{2} - \theta_H)^2] - c(\alpha^A) \\
&= -(1 - \alpha^A \tau) \frac{(\theta_H - \theta_L)^2}{4} - c(\alpha^A).
\end{aligned} \tag{15}$$

Given our assumptions about $c(\alpha)$, an interior solution is assured. The DM prefers the “Astroturf” equilibrium if $E(G^A) > G^{NA}$, that is, if

$$-(1 - \alpha^A \tau) \frac{(\theta_H - \theta_L)^2}{4} - c(\alpha^A) > -c(\alpha^{NA})/2. \tag{16}$$

It is evident that the “Astroturf” equilibrium is preferred by the DM if α^{NA} is very high and/or the audit cost function is highly convex. Thus, we have the following proposition:

Proposition 2 *An astroturf equilibrium exists when conditions (4), (9), (10), (14) and (16) hold.*

Recall that $\alpha^{NA} > \frac{1}{\tau} - \frac{l}{\tau\beta(\theta_H^2 - \theta_L^2)} = [\beta(\theta_H^2 - \theta_L^2) - l]/\tau\beta(\theta_H^2 - \theta_L^2)$. Thus, the size of τ is critical to determining which equilibrium can be supported. When τ is small, it becomes difficult for the DM to deter astroturf and α^{NA} becomes large. At the same time, in the “Astroturf” equilibrium, the marginal value of auditing declines so the DM audits less frequently. This increases the loss due to using a policy that doesn’t match the true state of the world, though it does decrease the DM’s expenditures on auditing. If $c(\alpha)$ is highly convex, then the increased costs of deterring astroturf will dominate, and the DM will be more likely to allow an astroturf equilibrium when τ is small.

In summary, we have demonstrated that a firm may be able to profitably engage in the practice of astroturfing, and that the DM may be unable to prevent this. Taken as a whole the results of this section lead to the following proposition.

Proposition 3 *The public decisionmaker would be better off if the firm were required to publicly disclose its expenditures on astroturf lobbying.*

Proof. Suppose the conditions in Proposition 2 hold. If public disclosure of expenditures on astroturf lobbying were required, then the DM would always be able to correctly infer the state, set the optimal policy for each state, and obtain expected payoff $G^0 = 0$. If the possibility of astroturf lobbying exists, one of two equilibria will result. In the “No Astroturf” equilibrium the DM’s expected payoff is $G^{NA} = -c(\alpha^{NA})/2 < G^0$, and in the “Astroturf” equilibrium the DM’s expected payoff is $G^A = -(1 - \alpha^A\tau)((\theta_H - \theta_L)^2/4) - c(\alpha^A) < G^0$. ■

Proposition 3 illustrates why decisionmakers may want to pass laws requiring the reporting of funding devoted to astroturf lobbying. Interestingly, this desire will exist even when efforts aimed at detecting astroturf are successful enough to deter the activity, since the DM must expend real resources on auditing to deter astroturf lobbying, and hence receives a strictly negative payoff even in the “No Astroturf” equilibrium.

A key feature of astroturf lobbying is its covert nature. Consequently, reporting provisions contained in early drafts of the 1995 Lobbying Disclosure Act would have eliminated this strategy by rendering it ineffective. A natural question to ask is whether such provisions would neutralize altogether corporate manipulation of the information provided provision by special interest groups. We find that this is not the case. In the following two sections we explore two alternative corporate strategies that also impede the ability of special interest group to provide information to the decisionmaker. The two strategies involve overt rather than covert actions on the part of the firm, and therefore would be unaffected by any public reporting requirements.

4 The “Bear Hug”

In this section we explore the use of publicly observable payments by the firm that are aimed at subsidizing the lobbying cost of special interest groups. We show that the firm may wish to make these payments to SIGs with either a negative or a positive bias. We focus on the case of a SIG with a positive bias, since our results are more striking, and perhaps counterintuitive, for this case. This case sheds light on the seemingly odd situation in which an interest group such as Greenpeace accepts funding from a large oil company such as Shell. That is, we assume, as in Section 2, that there exists a positive-biased SIG for which conditions (1) through (3) hold.²⁴ Then, as we have shown, there exists an equilibrium in which the SIG’s lobbying activity fully reveals to the DM the true state of the world. In this section we explore the firm’s relationship with the SIG under these circumstances. In particular, we consider a strategy in which the firm subsidizes the SIG’s lobbying activities,

²⁴All of our results in this section also go through if the SIG is negative-biased.

and examine when such a strategy might be profitable. This question is addressed in the following lemma.

Lemma 4 *If conditions (1) through (3) or (4) through (6) hold, then the firm has incentives to subsidize the SIG's lobbying activity if $l \leq l^f \equiv \beta(\theta_H - \theta_L)^2/4$.*

Proof. If the SIG's lobbying is informative, then the firm's expected payoff is $E(F) = -\beta\theta_L^2/2 - \beta\theta_H^2/2 = -\beta(\theta_L^2 + \theta_H^2)/2$. Alternatively, if the DM sets a policy simply based on its prior, the firm's payoff is $F = -\beta(\theta_H + \theta_L)^2/4$. Let the difference between these two payoffs be denoted by $\Delta = -\beta(\theta_H + \theta_L)^2/4 - [-\beta(\theta_L^2 + \theta_H^2)/2] = \beta(\theta_H - \theta_L)^2/4 > 0$. Thus, the firm is willing to spend up to this amount to subsidize the SIG's lobbying activity. ■

The intuition behind the lemma is straightforward. The firm's payoff is concave with respect to the stringency of policy. It faces very high costs from a policy of $p = \theta_H$, and thus has incentives to take action to avoid this outcome. By committing to subsidize the SIG, the firm effectively undermines the SIG's credibility with the DM (the SIG can no longer show the intensity of its preferences), and reduces the DM to adopting the policy $p = (\theta_L + \theta_H)/2$, its optimal choice when the state of the world is unknown. The strategy is thus a form of "signal jamming," similar in spirit to the analysis of Fudenberg and Tirole (1986) in the context of predation.²⁵

We use the term "bear hug" to refer to a strategy in which the firm embraces its opposition, clasping it so close as to smother it and reduce its effectiveness. The proposition shows that the firm can benefit from a policy of "bear hugging," that is, undertaking actions such as funding the collection or reporting of information by environmental or local community organizations.

Note that the firm's incentives to engage in a "bear hug" are proportional to $\beta(\theta_H - \theta_L)^2$. Hence, the value of this strategy grows as the gap between the high state and the low state grows. This is quite intuitive, since the bear hug can be seen as a form of insurance against costly stringent policies. Furthermore, *ceteris paribus*, less efficient firms (those with a high β) have more incentive to adopt this strategy than do more efficient firms.

Even when the firm wishes to offer the bear hug, the SIG must be willing to accept the firm's support. This will only be true if the SIG prefers to costlessly obtain the average policy outcome rather than incur the lobbying cost l in the high state to credibly deliver the

²⁵Note that the lobbying activities of a negatively biased SIG can also inform the DM of the true state; in this case, the SIG only lobbies when the state is low. Since the firm prefers that the DM not know the state of the world, signal jamming through the use of the bearhug can be valuable for a negatively biased SIG as well as one with a positive bias.

report θ_H . Mathematically, the SIG must prefer

$$.5[-((\theta_H + \theta_L)/2 - \theta_H - \delta)^2] + .5[-((\theta_H + \theta_L)/2 - \theta_L - \delta)^2] \quad (17)$$

to

$$-\delta^2 - l/2. \quad (18)$$

Expanding (17) and comparing it to (18) we see that the SIG is willing to accept the subsidy l if

$$l \geq l^{BH} \equiv (\theta_H - \theta_L)^2 / 2. \quad (19)$$

Comparing the condition (19) to the firm's subsidy condition $l \leq \beta (\theta_H - \theta_L)^2 / 4$, we are led immediately to the following proposition.

Proposition 5 *Assume conditions (1) through (3) or (4) through (6) hold. Then for $\beta > 2$ there exists a non-empty set of values $l \in [(\theta_H - \theta_L)^2 / 2, \beta (\theta_H - \theta_L)^2 / 4]$ for which the bear hug is profitable to the firm and will be accepted by the SIG in equilibrium.*

Figure 3 illustrates the existence of various bear hug equilibria. The curve l^{BH} provides the locus of lobbying costs l for which the SIG is willing to accept the bear hug over the relevant range of $a \equiv (\theta_H - \theta_L)$, i.e. $a < 2|\delta|$. The curve l^f illustrates, over the same range, the maximum subsidy the firm is willing to pay. Since l^f exceeds l^{BH} , there exist (a, l) pairs such that the SIG will be willing to accept the firm's bear hug. Note, however, that the firm will only engage in the bear hug for those values of a for which $l^{BH} > \underline{l}$. If this condition is violated the bear hug is not necessary, as the SIG's report lacks credibility.

As with the firm, we have examined the SIG's willingness to accept the firm's subsidy from an *ex ante* perspective. If the SIG knew that the true state of the world was high, the SIG would reject the firm's subsidy. Consequently, acceptance of the subsidy would reveal that the true state of the world was θ_L and the bear hug strategy would fail.²⁶

The effects of the bear hug on the DM's expected utility are shown in the following proposition.

Proposition 6 *The bear hug strategy reduces the public decisionmaker's expected payoff relative to the full information case.*

²⁶In the text we have considered only the existence of pure strategy equilibria. Mixed strategy equilibria in which the firm randomizes its subsidy offers are also possible, and may be more profitable for the firm. A proof is available from the authors upon request. Note that in a mixed strategy equilibrium, the DM does not observe directly whether the subsidy took place; rather, it simply believes (perhaps based on the firm's past behavior) that the firm is engaging in mixing behavior.

Proof. Without the subsidy, the DM’s expected utility is $E(G) = 0$. The SIG can be relied upon to reveal the true state, and the DM can thus tailor policy perfectly to each state of the world. When the firm subsidizes the SIG, the DM’s expected utility is $E(G^{BH}) = (1/2)[-((\theta_L + \theta_H)/2 - \theta_L)]^2 + (1/2)[-((\theta_L + \theta_H)/2 - \theta_H)]^2 = -(\theta_H - \theta_L)^2/4 < 0$. Hence the DM is worse off when the firm supports the SIG. ■

The proposition shows that under conditions (1) through (3) or (4) through (6), the DM is strictly worse off when the firm provides financial support to the SIG. While signal-jamming is profitable for the firm, and may be accepted by the SIG as a way to economize on lobbying costs, it is unwelcome to the decisionmaker because it prevents the optimal matching of policy to circumstances.

There are two potential issues in assessing when the “bear hug” is a viable strategy. First, the strategy must apply to situations where the true state of the world is unknown to all players at the time the subsidy is granted. The reason for this restriction is as follows. If the firm knew the true state of the world was θ_L , it would prefer that the conditions of truthful revelation held. These conditions would require that no subsidy be given, so the SIG eschews lobbying. If the firm knew the state was θ_H , however, it would want to publicly make a subsidy payment to the SIG so as to undermine its credibility before the DM. Thus, if the firm knew the true state, then its subsidy would be state dependent, and the DM could determine the true state simply by observing whether the subsidy payment had been made. In consequence, the bear hug strategy is more likely to apply to situations with true scientific uncertainty or situations with a risk of accidents than to situations where the firm knows the state in advance. The bear hug can thus be seen as a kind of insurance policy against worst-case policy outcomes.

The second issue affecting the viability of the bear hug is that the firm must be able to ensure that its subsidy is used to subsidize the SIG’s lobbying costs on the particular issue of concern. Thus, there may be difficulties implementing the bear hug strategy if the SIG operates in multiple policy arenas. Returning to our example, a general purpose gift to an environmental group may simply go to subsidize the group’s fixed costs, but may not guarantee that extra funds are devoted to lobbying about dioxin. Thus, the firm may need to tie the gift to SIG activity in a particular issue area. This might be done by providing the SIG with a forum in which it can express its views. For example, in the paper industry example, the environmental group could be invited to participate in a paper industry forum, at the industry’s expense, thereby targeting the support toward a particular issue.

5 Self-Regulation

In section 3 we illustrated that the *ex post* strategy of astroturf is only possible with a SIG that has a negative bias. In section 4 we analyzed the bearhug, which can be used with either type of SIG. This strategy, however, is dependent on the cooperation of the SIG in that it must accept the firm's payment *ex ante*, and credibly commit to use the payment to subsidize its lobbying costs on the issue at hand. In this section, we study an *ex ante* corporate strategy that does not require any cooperative actions by the SIG. Specifically we study the possibility that the firm may be able to alter the SIG's lobbying behavior by reducing the severity of the high state, *i.e.* to reduce θ_H , through voluntary improvements made *ex ante*. This might be done, for example, through design measures for a new facility that reduce the impact of worst-case outcomes. As long as the DM has the power to hold the firm to the design it proposes, such actions constitute credible commitments.

The basic intuition here is that if the difference between the high and low states is sufficiently small, then the SIG will have little motivation to lobby the DM. Hence, self-regulation by the firm may induce the SIG to eschew lobbying, with the result that the DM sets the average policy. However, there may also be a counterintuitive reason for self-regulation: as we show below, the firm may also obtain the average policy because self-regulation induces the SIG to lobby in both states. In either case, the firm's profits rise by $\Delta = \beta(\theta_H - \theta_L)^2/4$, as shown in Lemma 2. Baron (2001) refers to such profit-driven self-regulation as "strategic corporate social responsibility," in contrast to corporate social responsibility that is altruistically motivated.

Recall the notation $a = \theta_H - \theta_L$, and denote by a_0 the initial gap between the states. We consider self-regulation as a voluntary reduction in a on the part of the firm, cutting it from a_0 to a_1 . That is, the firm's voluntary action reduces the severity of the DM's optimal policy in the high state of the world. Thus, if the firm reduces θ_H from $\theta_L + a_0$ to $\theta_L + a_1$, and this induces the SIG to eschew lobbying or to lobby in both states, then the DM sets the average policy, *i.e.* $p = \theta_L + a_1/2$, and the firm's payoff is $F^{SR} = -\beta(\theta_L + a_1/2)^2 = -\beta(\theta_L^2 + a_1\theta_L + a_1^2/4)$.²⁷ If the firm took no action, and the SIG revealed the true state through its lobbying decisions, then the firm's expected payoff would be $F^0 = -\beta(\theta_H^2 + \theta_L^2)/2 = -\beta(\theta_L^2 + a_0\theta_L + a_0^2/2)$. The net benefit to the firm is $\Delta^{SR}(a_1) = F^{SR} - F^0 = \beta(a_0 - a_1)\theta_L + \beta(2a_0^2 - a_1^2)/4 > 0$.

Recall that the payoff function for the firm is $F = -\beta p^2$. How should we represent the cost of achieving a ? If the firm were forced to comply with a policy of $p = \theta_H$, the cost difference between a_0 and a_1 would be $k(a_1) = -\beta(\theta_L + a_1)^2 - (-\beta(\theta_L + a_0)^2) = \beta(2\theta_L + a_0 + a_1)(a_0 - a_1)$.

²⁷We could allow the firm to reduce both θ_H and θ_L , and as long as the former is reduced more than the latter, all our results in this section would still go through.

In order to be consistent with this payoff function, we will assume that the cost to the firm of achieving such a reduction is $k(a_1) = \beta(2\theta_L + a_0 + a_1)(a_0 - a_1)$. Thus we allow the firm no extra benefit from engaging in unmandated reductions in θ_H .²⁸

Combining the benefits and the cost of voluntary action (assuming the action is sufficient to render the SIG's lobbying choice uninformative) gives a net payoff to voluntary action of

$$V(a_1) = \Delta^{SR}(a_1) - k(a_1) = \frac{3}{4}\beta a_1^2 - \frac{1}{2}\beta a_0^2 - \beta\theta_L(a_0 - a_1).$$

Recalling that $a_1 < a_0$, it is easy to see that $V(a_1)$ is positive for a_1 sufficiently close to a_0 . Since $\partial V/\partial a_1 > 0$, the firm prefers the largest a_1 (smallest amount of self-regulation) that is sufficient to render the SIG's lobbying choice uninformative.²⁹ Let \bar{a}_1 be smallest value of a_1 that the firm is willing to choose; this can be found by setting $V(\bar{a}_1) = 0$. Thus, the firm is willing to self-regulate to any value of a between \bar{a}_1 and a_0 if this will undermine the SIG's lobbying efforts.³⁰

How will a voluntary action affect the decisions of the SIG? Recall that conditions (2) and (3) define when the SIG will find it worthwhile to lobby before the DM. If $l \leq \bar{l} = a(2\delta + a)$ then the SIG finds it worthwhile to lobby when the state is θ_H , while if $l \geq \underline{l} = a(2\delta - a)$ then the SIG does not lobby when the state is θ_L . Clearly self-regulation shifts \bar{l} and \underline{l} , and in the process may cause the SIG to change its lobbying behavior. This is perhaps most easily seen by reference to Figure 1, in which \bar{l} and \underline{l} divide the (a, l) space into three regions: 1) The region with $l > \bar{l}$, in which the SIG never lobbies, 2) the region with $l < \underline{l}$, in which the SIG always lobbies, and 3) the region with $\bar{l} > l > \underline{l}$, in which the SIG lobbies only if $\theta = \theta_H$. Of course, only in the third region is lobbying activity actually informative to the DM. Figure 4 builds on Figure 1 but adds two shaded regions. In each of these regions, the initial point (a_0, l) is within the region where the SIG's lobbying efforts are informative, but the firm is willing to self-regulate to an extent that will result in a westward move that causes the SIG to change behavior. In the shaded region close to \bar{l} , self-regulation causes the SIG to abandon lobbying. In the shaded region close to \underline{l} , self-regulation has the opposite effect: it induces the SIG to lobby in both states of the world. In either case, however, the SIG's lobbying choice becomes uninformative for the DM, which is profitable for the firm.

²⁸It is not uncommon in the literature on voluntary environmental agreements for authors to assume that voluntary actions are less costly than mandated actions. The authors of these papers argue that voluntary actions allow firms greater flexibility in meeting environmental goals. We refrain from modeling this exogenous bias, which serves only to make voluntary actions more desirable. For a discussion of papers that adopt this exogenous cost bias in favor of voluntary actions see Lyon and Maxwell (2002b).

²⁹Recall that the SIG's lobbying can be made uninformative either by inducing the SIG to never lobby, or to lobby in both states of the world.

³⁰Note that $\bar{a}_1 = -\frac{2\beta}{3}\theta_L + \frac{2\beta}{3}\sqrt{(\theta_L^2 + 3\theta_L a_0 + 3a_0^2/2)}$, and is thus a function of a_0 . However, in what follows we suppress this dependence for notational simplicity.

We summarize these results in the following proposition.

Proposition 7 *There exist two ways in which corporate self-regulation can profitably alter interest group lobbying behavior: a) if $a_0(2|\delta| + a_0) > l > a_0(2|\delta| - a_0)$ and $l > (a_0 - \bar{a}_1)(2|\delta| + a_0 - \bar{a}_1)$, then self-regulation induces the interest group to eschew lobbying, and b) if $a_0(2|\delta| - a_0) < l < (a_0 - \bar{a}_1)(2|\delta| - a_0 + \bar{a}_1)$, then self-regulation induces the interest group to become a pure advocate that lobbies regardless of the actual state of the world. In either case, the interest group's lobbying behavior becomes uninformative and profits rise.*

Proof. Begin with case (a). We require two conditions. First, $a_0(2|\delta| + a_0) > l > a_0(2|\delta| - a_0)$ ensures that the initial pair (a_0, l) is such that the SIG's lobbying behavior is informative; that is, it ensures that $l \in (\underline{l}, \bar{l})$. Second, $l > (a_0 - \bar{a}_1)(2|\delta| + a_0 - \bar{a}_1)$ ensures that after self-regulation, the SIG is in the region in which lobbying is never worthwhile; that is, after self-regulation, we have $l > \bar{l}$. Now turn to case (b). The condition $a_0(2|\delta| - a_0) < l$ ensures that the SIG does not always lobby at the initial pair (a_0, l) . The second condition $l < (a_0 - \bar{a}_1)(2\delta - a_0 + \bar{a}_1)$ ensures that after self-regulation, we have $l < \underline{l}$, and the SIG engages in (uninformative) lobbying regardless of the state of the world. ■

The effect of self-regulation on interest group behavior in this model is more complex than in some previous models. For example, Maxwell, Lyon and Hackett (2000) show that self-regulation can deter entry into the political arena by a consumer interest group. Self-regulation reduces the marginal benefits of political entry, and thus can effectively deter it. In the present model, this effect corresponds to case (a) of the above proposition. However, the present model also admits case (b), in which self-regulation induces *greater* lobbying activity on the part of the interest group. Why does this occur? Recall that the SIG has a positive bias, meaning it prefers a policy p that is strictly greater than the true state. Furthermore, we have focused on the case where the SIG's bias is great enough ($\delta > a/2$) that the SIG has an incentive to report falsely when $\theta = \theta_L$. If it can thereby induce the DM to set a policy of $p = \theta_H$, it obtains payoff $-(a - \delta)^2 - l$ rather than the payoff of $-\delta^2$ it would receive for truthful reporting that yields a policy of $p = \theta_L$. Thus, the payoff from misreporting rises as a approaches δ from above. Self-regulation thus makes false reporting more advantageous for the SIG, and leads the SIG to lobby more often.

It is also interesting to ask how the decisionmaker is affected by self regulation that renders lobbying uninformative. This is the subject of the following proposition.

Proposition 8 *Self-regulation reduces the public decisionmaker's expected payoff relative to the case of full information.*

Proof. Without self regulation, the DM’s expected payoff is simply $G^0 = 0$, since policy can be tailored precisely to the state of the world *ex post*. With self regulation, the DM lacks information about the state and must set the average policy. Suppose that the firm’s self-regulatory actions are observed by the DM. Then the DM sets $p = \theta_L + a_1/2$, which is equal to the expected value of the state. Now the DM’s expected payoff becomes

$$E(G^{SR}) = -.5(\theta_L + a_1/2 - \theta_L)^2 - .5(\theta_L + a_1/2 - \theta_H)^2 = -a_1^2/2 < 0.$$

Clearly $E(G^{SR}) < G^0$, and the DM is worse off as a result of the firm’s self-regulatory action. ■

In this model, the DM is best able to maximize his objective function when he has full information about the state of the world. Then he can tailor policy to the specifics of the situation before him. Self-regulation is only undertaken by the firm if it will render the SIG’s lobbying uninformative. This deprives the DM of the information he desires, and as a result the DM is worse off. This result contrasts with that in earlier work, such as that of Lyon and Maxwell (2002a), who show that the regulator benefits when industry self-regulation preempts the imposition of new regulations. The key difference is that Lyon and Maxwell (2002a) study a model in which self-regulation does not affect the information flow to the regulator.

It is worth noting that if the firm undertakes a strategy that renders lobbying uninformative, then the DM obtains higher utility from self-regulation than from the bear hug.³¹ Nevertheless, although self-regulation reduces the severity of the high state, it never entirely compensates for the loss of information caused by the decision to self-regulate.

6 Multiple Interest Groups

To this point, we have concentrated on cases involving only a single interest group. In this section, we discuss how our results may be extended to cases with multiple interest groups. We follow the typology used by Grossman and Helpman (2001) to classify the structure of multiple SIG situations: 1) “Like bias” arises when all groups share the same direction of bias, but with different intensity; 2) “Opposite bias” arises when different groups are biased in opposite directions, and 3) “Unknown bias” arises when the groups receive imperfect signals regarding the state of the world. The first two cases, in contrast to the third, assume

³¹We do not formally compare the DM’s expected payoff in the astroturf equilibrium with that obtained under the bearhug or self-regulation. This is because the astroturf strategy is only applicable when the firm already knows the state of the world, while the other two strategies are only applicable when the firm does not know the state.

that both SIGs have perfect information regarding the state of the world at the time they lobby the DM. We consider these in turn, focusing on the case of two SIGs for simplicity.

6.1 Like Bias

We will label the two SIGs “radical” and “moderate,” with the former possessing a larger value of δ . We assume the moderate group, as in earlier sections, has a bias that satisfies conditions (1) through (3). The more radical group may meet these conditions, but could also be so biased that it always lobbies and always claims that the state is high. This latter possibility may arise even if lobbying is costly, if the radical group’s bias (δ) is high enough. If the firm prefers a policy set at the average level, then it prefers to mute (render uninformative the group’s lobbying actions) the moderate group, since the radical group lacks credibility anyway. This can be accomplished by bear hugging the moderate group *ex ante*, if that group’s bias is great enough that it will always claim the state is high when lobbying is costless. (Alternatively, similar results can be achieved through astroturfing *ex post*, if the group has a negative bias.) Thus, this case differs little from the single SIG case analyzed above.

Alternatively, if the radical SIG’s bias is not too great, then the DM could also rely on it to provide reliable information through costly lobbying. In this case, bear hugging (or astroturfing) the moderate SIG will not be sufficient to affect the DM’s decision. Instead, the firm must subsidize both SIGs. Again, however, this case differs only trivially from the case of a single SIG.³²

Overall, we conclude that the addition of a second SIG with like bias to that of the first SIG is unlikely to generate much additional insight. However, it is worth noting that if all groups must be subsidized, then the cost of any kind of subsidy strategy rises linearly with the number of SIGs. This is not true of the self-regulation strategy, however. A single voluntary improvement affects all SIGs at once. If the firm undertakes enough voluntary action to preempt the involvement of the most extreme group, then all other groups will be preempted as well. Thus, we hypothesize that self-regulation is likely to outperform subsidy strategies as the number of SIGs grows.³³

³²Equilibria in all models with incomplete information depend critically on the beliefs held by the players. In our model, the actions of the firm depend on how the DM chooses to interpret the lobbying actions of the SIGs. For example, suppose the DM held the belief that the state is high if both SIGs lobby and is low if neither firm lobbies, and he views any other outcome as uninformative. Then the firm would achieve the average policy by muting only one SIG.

³³This is particularly likely to be the case if the bias of the most radical group does not change as the number of groups grows. Otherwise, if the bias of the most radical group grows along with the number of groups, then the cost of self-regulation will grow as well, and the cost comparison becomes more difficult.

6.2 Opposite Bias

When the two SIGs are biased in opposite directions, matters become more interesting. At least two types of equilibria are possible: 1) the DM ignores one SIG and simply relies on the other, and 2) Each SIG lobbies in one state of the world, and the DM relies on both. In particular, an equilibrium of the second type may exist in which the SIG with positive bias lobbies in the high state, while the SIG with negative bias lobbies in the low state.³⁴

Recall from our earlier analysis that the bear hug can be applied to groups with either type of bias, but requires commitment ability and must be undertaken before the firm learns the state of the world. Astroturf does not require commitment ability, and can be undertaken *ex post*, but it can only be employed with groups having a negative bias. Self-regulation is undertaken *ex ante*, and will influence both types at once.

Case 1 is similar to the case of like bias. If the firm successfully bear hugs the “active” SIG, then the inactive SIG may find it worthwhile to lobby, and the DM will find it worthwhile to pay attention to it. Thus, the firm needs to bear hug both of the SIGs. Alternatively, the firm may use self-regulation to influence both SIGs at once.

Case 2 is somewhat more complex. On the one hand, if the lobbying activities of one group are rendered ineffective, then the initial equilibrium is destroyed. However, there is an alternative equilibrium (Case 1) in which the DM pays attention to only one of the SIGs, and this becomes the only equilibrium if one SIG is bear hugged. Hence, the firm must again undermine the credibility of both groups, either through bear hugs or self-regulation, if it wishes to be successful. On the other hand, if the firm knows the state of the world, it might choose to engage in astroturf lobbying when the state is high, thereby inducing the negatively biased SIG to lobby at the same time that the positively biased SIG lobbies. The effectiveness of this strategy will depend on the DM’s beliefs in this out-of-equilibrium event. If, as seems reasonable, the DM sets the average policy when both SIGs lobby, then astroturf may be profitable in exactly the same way as in section 3 above.

The general point is that the basic structure of our analysis appears to remain valid in the presence of multiple SIGs, as long as those SIGs all possess full information regarding the state of the world. The main change from adding multiple groups is that self-regulation may become relatively more attractive as the number of SIGs rises.

6.3 Unknown Bias

As before, one group is assumed to be radical, and willing to lobby in both states of the world. However, the DM is assumed to be unable to distinguish one group from the other,

³⁴See Grossman and Helpman (2001), chapter 5, for details.

hence can only make use of information regarding the *number* of firms that lobby. Lohmann (1993) analyzes this setting in the context of $N > 2$ groups, but Grossman and Helpman (2001) show that her main insights can be derived in a model with just two groups. Consider the case of two groups with like biases. Lohmann emphasizes the case in which the more radical SIG always lobbies, regardless of the state. The DM does not know which group is the more biased, but can still use the extent of lobbying as a noisy signal regarding the state of the world. For example, the DM may conclude that the state is high if two SIGs lobby, and low if only one does.³⁵ If the firm can identify the more moderate SIG, then it can subsidize the moderate group, just as in the case of known bias. If this is not possible, then the firm must subsidize both groups.

Now consider the case of opposite bias. Suppose that the SIG with positive bias is the more radical one, and it plays the role of a pure advocate, that is, it always lobbies and claims the state is high. The more moderate SIG only lobbies when the state is low. Thus, the appearance of 1 SIG indicates that the state is high, while the appearance of 2 groups indicates the state is low, and the DM sets a low level of policy when both groups lobby, but a high level when only one group lobbies. Once again, if the firm subsidizes the moderate group, then that group will always lobby, and the DM must set policy without gaining any information from the SIGs. If the firm cannot determine which group is which, then it must subsidize both.

The case of unknown biases is more subtle than the first two cases we discussed, since the SIGs don't know the state of the world for certain. Sometimes they will be wrong, and the DM must take this into account. Nevertheless, for our purposes, most of the qualitative features of the models seem basically the same.

One new possibility may emerge in the case of $N > 2$ SIGs with imperfect information. If the groups move sequentially in presenting their information, the possibility of information cascades arises, as in Bikhchandani, Hirshleifer, and Welch (1992). In such a setting, there is a large premium to being the first SIG to lobby, since all the subsequent SIGs may be influenced by the actions of the first. There is also a large premium to the firm if it can influence the information revealed by the first SIG to lobby. This case, while interesting, is beyond the scope of the present paper.

³⁵A failure to lobby by both firms is off-equilibrium path behavior. Grossman and Helpman (2001, p. 154) identify beliefs for the DM under which it infers the state is low when neither firm lobbies.

7 Conclusions

In this paper, we have developed a model to explore how firms may influence the lobbying behavior of special interest groups. We built on the framework presented by Grossman and Helpman (2001), in which costly lobbying may convey unverifiable information to a public decisionmaker. The basic idea of this framework is that when lobbying is costly, an interest group’s decision to lobby provides credible information about the strength of its preferences regarding a particular policy issue. We have shown that firms may have both the incentive and the ability to undermine this information transmission process, reducing the public decisionmaker’s payoff in the process.

We considered three corporate non-market strategies: 1) “astroturf ” in which the firm subsidizes the lobbying activities of a group with similar views, 2) the “bear hug,” in which the firm subsidizes the lobbying activities of an interest group, and 3) self-regulation, in which the firm voluntarily limits the potential social harms from its activities. All three of these strategies can be used to reduce the informativeness of lobbying, which can be profitable for the firm if the costs of complying with public policy are sufficiently convex. When compliance costs are convex, the firm gains if the public decisionmaker sets policy based on expected or average social harm, rather than face the risk that policy will be tailored to actual harm.

In many situations, the firm is likely to know the true state of the world already, especially if that state depends on characteristics of the firm’s technology or management processes. For example, the state of the world might be the level of health risk associated with the operation of a particular plant, which depends upon corporate decisions regarding technology and management. In such settings, astroturf lobbying can be induced by the firm, which covertly subsidizes the lobbying activity of an interest group with similar preferences in states of the world where the interest group would not otherwise lobby. For example, the group might represent local business organizations that stand to benefit if the firm builds a new plant in the area. We model this strategy as a form of costly state falsification. We show that the decisionmaker has incentives to audit the relationship between the firm and the interest group for evidence of astroturf lobbying, and identify conditions under which astroturf lobbying nevertheless takes place in equilibrium. Our model shows that a law requiring the reporting of astroturf lobbying expenditures would render the strategy ineffective, and that this would be desirable for the public decisionmaker.

Requiring the reporting of astroturf lobbying expenditures is worthwhile, but is not a panacea. We examine two alternative corporate strategies that can also reduce the informativeness of lobbying, even when their use is common knowledge to all players. These strategies differ from astroturf lobbying in that they can only be used by the firm before it

learns the true state of the world. This is particularly likely in situations of true scientific uncertainty, such as currently exists regarding the future impacts of global warming. These alternative strategies can prevent special interest groups from informing the decisionmaker about the true state after they learn its value.

The “bear hug” serves as a signal-jamming device that prevents the interest group from signalling the intensity of its views. One might expect that the group would be unwilling to accept a subsidy that reduces the credibility of its statements. Nevertheless, we show that if lobbying is costly enough, then it is optimal for the group to accept the firm’s embrace. It is important to note that this strategy may not be dynamically consistent for the firm: even though it raises expected profits *ex ante*, it is unprofitable *ex post* in some states of the world. Hence, the strategy is only feasible if the firm can credibly commit to subsidize the interest group regardless of the true state.

The third strategy we study is self-regulation, namely, voluntary actions to reduce the social harm that occurs in adverse states of the world. Such voluntary actions can change the lobbying incentives of interest groups, and may render them uninformative, which is profitable for the firm. Self-regulation has subtle effects in our model. The most intuitive effect is that self-regulation can preempt interest group lobbying, by reducing the benefit from lobbying relative to its cost. Another, less intuitive, possibility is that self-regulation can strengthen the incentives of a positive-biased interest group to falsely report that the state is high when it is really low. An interest group with positive bias wants a policy greater than that justified by the true (low) state of the world, but it may not want the policy to be fully as stringent as would be justified in the high state of the world. By bringing the high state closer to the low state, self-regulation makes it less costly for the interest group to endure the stringent policy, and makes it more attractive for the group to engage in lobbying in both states of the world.

Our analysis focuses on the case of a single interest group, but appears to be robust to the incorporation of multiple groups. The most interesting possibility that arises with multiple groups is that self-regulation becomes relatively more attractive, since a single investment in self-regulation can mute all groups at once, while the cost of a strategy based on subsidies rises linearly with the number of groups.

Under all three of the strategies we consider, the public decisionmaker is made worse off. The key reason is that when the decisionmaker is fully informed, he can tailor policy precisely *ex post* to the particular state of the world. All three of the strategies we study here are designed to stem the flow of information, and while this increases profits it simultaneously reduces the decisionmaker’s expected payoff.

References

- [1] Aumman, R.J. and A. Brandenburger. 1995. "Epistemic Conditions for Nash Equilibrium," *Econometrica*, 63: 1161-1180.
- [2] Baron, David P. 2001. "Private Politics, Corporate Social Responsibility, and Integrated Strategy," *Journal of Economics and Management Strategy*, 10: 7-46.
- [3] Beckenstein, Alan R., Brad Webb, Frederick J. Long, and Barbra L. Marcus. 1994. "Chlorine and the Paper Industry," in *Stakeholder Negotiations: Exercises in Sustainable Development*, edited by Alan R. Beckenstein, Frederick J. Long, Matthew B. Arnold and Thomas N. Gladwin, 1995, Richard D. Irwin.
- [4] Becker, Gary S. 1983 "A Theory of Competition among Pressure Groups for Political Influence," *Quarterly Journal of Economics*, 98: 371-400.
- [5] Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch. 1992. "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy* 100: 992-1026.
- [6] Craig, Tim. 2002. "Astroturf Lobbying: Community Leaders Decry Lobby Firm's Fax," *The Baltimore Sun*, March 9, 2002. <http://www.commondreams.org/headlines02/0309-04.htm>
- [7] Crawford, Vincent and Joel Sobel. 1982. "Strategic Information Transmission," *Econometrica* 50: 1431-1451.
- [8] Crocker, Keith J. and Sharon Temmyson. 1997. "Contracting with Costly State Falsification: Theory and Empirical Results from Automobile Insurance," mimeo, University of Michigan Business School.
- [9] DeSimone, Livio D. and Frank Popoff. 2000. *Eco-Efficiency: The Business Link to Sustainable Development*. Cambridge, MA: The MIT Press.
- [10] Fudenberg, Drew and Jean Tirole. 1986. "A 'Signal-Jamming' Theory of Predation," *The Rand Journal of Economics*, 17: 366-376.
- [11] Greer, Marsha L. and Anna K. Harding. 1993. "The Health Impact of Hazardous Waste Sites on Minority Communities: Implications for Public Health and Environmental Health Professionals." *Journal of Environmental Health* 55(77): 6-9.
- [12] Grossman, Gene M. and Elhanan Helpman. 2001. *Special Interest Politics*. Cambridge, MA: The MIT Press.
- [13] Krishna, Vijay and John Morgan. 2001. "A Theory of Expertise," *Quarterly Journal of Economics*, 116: 747-775.

- [14] Lohmann, Suzanne. 1993. "A Signaling Model of Informative and Manipulative Political Action," *American Political Science Review*, 87: 319-333.
- [15] Lyon, Thomas P. and John W. Maxwell. 2002a. "Self-Regulation, Taxation, and Public Voluntary Environmental Agreements," *Journal of Public Economics*, forthcoming.
- [16] Lyon, Thomas P. and John W. Maxwell. 2002b. "Voluntary Approaches to Pollution Abatement: A Survey," in Maurizio Franzini and Antonio Nicita, eds., *Economic Institutions and Environmental Policy*, Ashgate Publishing.
- [17] Maxwell, John W., Thomas P. Lyon, and Steven C. Hackett. 2000. "Self-Regulation and Social Welfare: The Political Economy of Corporate Environmentalism," *Journal of Law and Economics*. 43: 583-617.
- [18] Milgrom, Paul and John Roberts. 1986. "Relying on the Information of Interested Parties," *RAND Journal of Economics*, 17: 18-32.
- [19] Olson, Mancur. 1965. *The Logic of Collective Action*. Cambridge, MA: Harvard University Press.
- [20] Peltzman, Sam. 1976. "Toward a More General Theory of Regulation," *Journal of Law and Economics*, 19: 211-248.
- [21] Sanchez, Samantha. 1996. "How the West Is Won: Astroturf Lobbying and the 'Wise Use' Movement," *The American Prospect* 25: 37-42.
- [22] Stauber, John and Sheldon Rampton. 1995. *Toxic Sludge is Good for You: Lies, Damn Lies and the Public Relations Industry*. Monroe, Maine: Common Courage Press.
- [23] Stigler, George J. 1971. "The Economic Theory of Regulation," *Bell Journal of Economics and Management Science*, 2: 1-21.
- [24] Taylor, D. 1992. "The Environmental Justice Movement." *EPA Journal* 18(1): 24-36.

Figure 1: Interest Group Lobbying Behavior




Figure 2: Feasibility of Astroturf Equilibrium




Figure 3: Incentives for the Use of the Bearhug

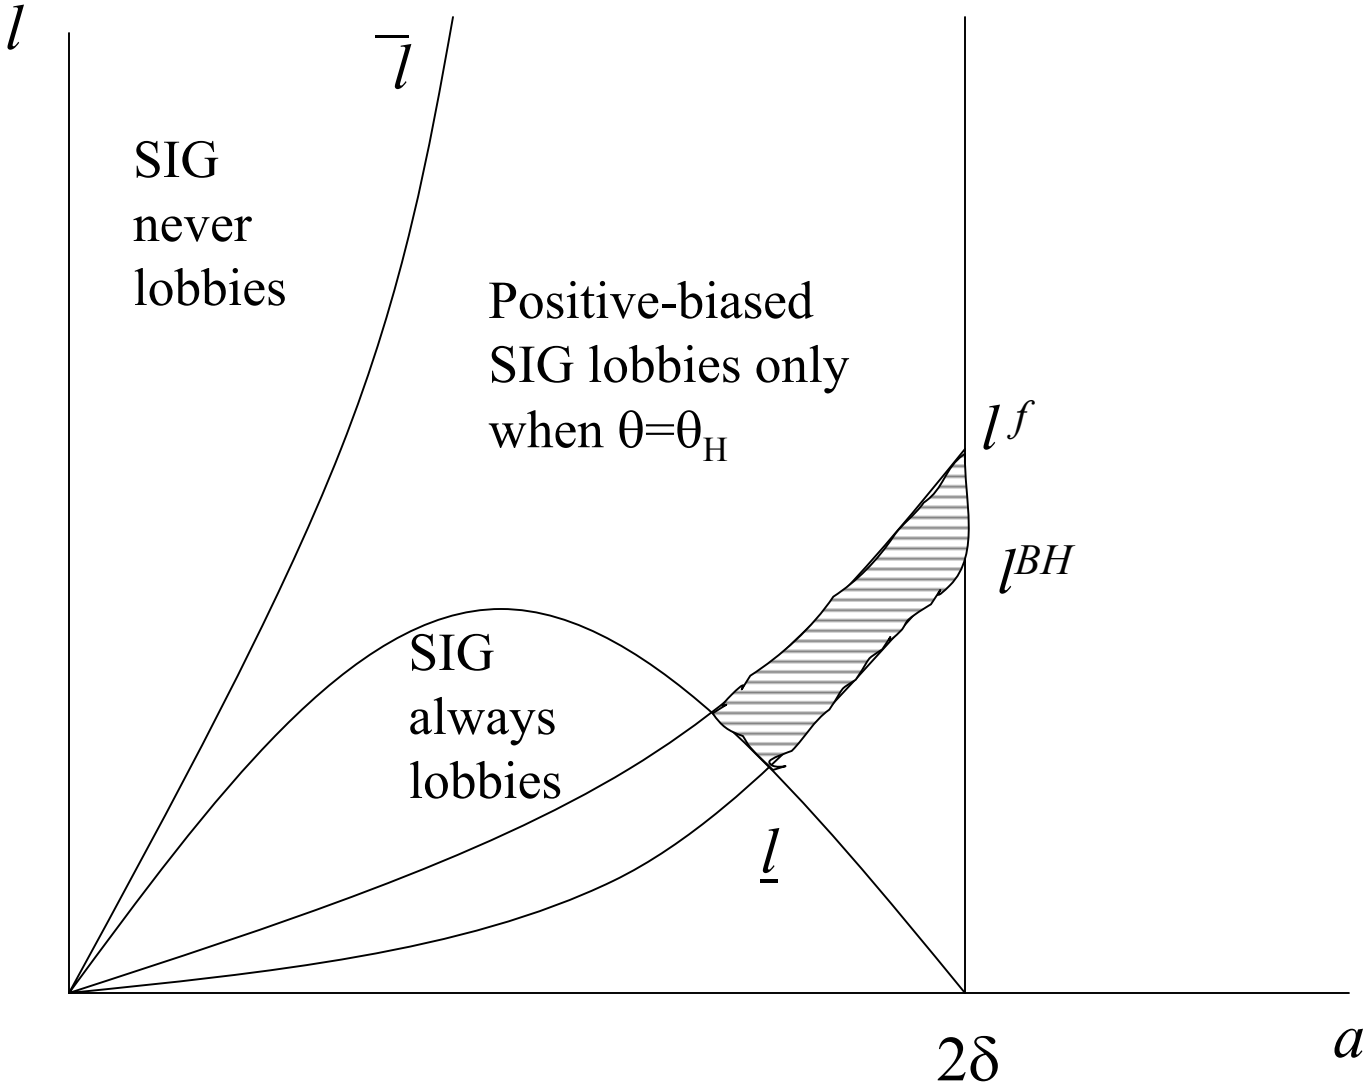


Figure 4: Profitability of Self-Regulation

