



HAL
open science

Merchants of doubt: Corporate political action when NGO credibility is uncertain

Mireille Chiroleu-Assouline, Thomas P Lyon

► **To cite this version:**

Mireille Chiroleu-Assouline, Thomas P Lyon. Merchants of doubt: Corporate political action when NGO credibility is uncertain. *Journal of Economics and Management Strategy*, 2020, 29 (2), pp.439-461. 10.1111/jems.12338 . halshs-02552465

HAL Id: halshs-02552465

<https://shs.hal.science/halshs-02552465v1>

Submitted on 23 Apr 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Merchants of Doubt: Corporate Political Action

when NGO Credibility is Uncertain*

Mireille Chiroleu-Assouline[†] and Thomas P. Lyon[‡]

June 25, 2019

Keywords: Informational lobbying, persuasion, nonmarket strategy, special interest politics

JEL Codes: D72, D82, L31, Q58

Suggested Running Head: Merchants of Doubt

*This research has been funded by the French National Research Agency (ANR-15-CE05-0008-01) and the program Investissements d'Avenir of the French government (ANR-10-LABX-93-01). We thank seminar audiences at the Canadian Resource and Environmental Economics workshop, Columbia University, the European Association of Environmental and Resource Economists, the Ivey Business School, the Paris Workshop on the Politicization of Firms, the Paris Workshop on Environmental NGOs and Public Policy, and the University of Michigan for their helpful comments.

[†]Paris School of Economics, University of Paris 1 Panthéon-Sorbonne and INRA UMR Economie Publique, Paris, France.

Email : Mireille.Chiroleu-Assouline@univ-paris1.fr.

[‡]Corresponding author. Ross School of Business, 701 Tappan St., University of Michigan, Ann Arbor, MI, 48109. Email: tplyon@umich.edu.

Abstract

The literature on special interest groups emphasizes two main influence channels: campaign contributions and informational lobbying. We introduce a third channel: providing information about the credibility of political rivals. In particular, non-governmental organizations (NGOs) often aim to communicate scientific knowledge to policymakers, but industry-backed groups often attempt to undermine their credibility. We extend a standard signaling model of interest-group lobbying to include fixed costs of policymaker action, and show that these costs make possible two mechanisms for creating doubt about the value of policy action. The first uses Bayesian persuasion to suggest the NGO may be a non-credible radical. The second involves creating an opposition think tank that acts as a possible radical, not a credible moderate. We show that the think tank implements the Bayesian persuasion benchmark, and we characterize how optimal think tank design varies with exogenous parameters.

1 Introduction

Political action to regulate externalities, provide public goods and protect public health depends crucially on scientific knowledge. However, public doubts about the credibility of scientific experts may undermine political support for regulation. Indeed, Oreskes and Conway (2010) present a series of in-depth case studies showing how business groups whose profits were threatened by regulation sowed doubt about the scientific basis for action. They cover cases involving the regulation of tobacco, acid rain, the hole in the ozone layer, and climate change. Their central contribution is to detail how business interests attempted to block or delay legislation, even though the weight of scientific evidence was against them. In each case, the strategies involved creating uncertainty about the scientific case for regulation by undermining the credibility of scientists and science-based non-governmental organizations (NGOs). In fact, a 1969 memo by a tobacco industry executive explained that “Doubt is our product since it is the best means of competing with the ‘body of fact’ that exists in the minds of the general public.” (Oreskes and Conway, 2010, p. 34).

One common strategy for creating doubt is to attack the character of science-based organizations, accusing them of a political bias that makes their recommendations untrustworthy. Even if there is no evidence to support such an attack, bold assertions may be enough to create sufficient doubt to damage a scientist’s reputation. One victim of such attacks was Ben Santer, recipient of a 1998 MacArthur Foundation “genius” grant and lead author of Chapter 8 of the Second Assessment Report issued by the Intergovernmental Panel on Climate Change (IPCC). Santer was the target of a 1996 Op-Ed piece in the *Wall Street Journal* that accused him of making changes to the IPCC report to “deceive policy makers and the public.” (Oreskes and Conway, 2010, p. 3). The charges were false—he had simply edited the chapter in response to review comments from fellow scientists—but that did not stop Frederick Seitz, chairman of the George C. Marshall Institute, from making them. Seitz, then 85 years old, had had an illustrious career in materials science and solid state physics. Upon his retirement in 1979, he became a consultant for the tobacco industry and later the fossil fuel industry, arguing against regulation of second-hand smoke and global warming. Although Santer and 40 other climate scientists wrote a sharp rebuttal to the *Journal*, the newspaper edited the

rebuttal heavily and deleted the names of the other 40 signatories, making it appear to be a self-serving response by Santer alone. The ensuing public war of words was too complicated for most members of the public to parse, and could easily be seen as just another debate between “two sides,” each with their own equally valid views. Doubt had been sown.

A second common strategy is for a “think tank” to put forward competing scientific “experts” whose opinions are opposed to those of the scientists calling for policy action. This is a tactic that has become increasingly common since Anthony Fisher created the Atlas Network in 1981 with the intention “to litter the world with free-market think tanks.” (Djelic, 2014).¹ For example, in 1989 the Marshall Institute issued a “white paper” on global warming that argued any warming was caused by solar activity, not greenhouse gas emissions. Although the paper was not written by climate scientists and had not been subjected to peer review, the Institute contacted the White House and was invited to present the report to members of the Council of Economic Advisers, the Office of Management and Budget, and other high-level executive branch entities. This outreach effort was remarkably effective and “had a big impact, stopping the positive momentum that had been building in the Bush administration.” (Oreskes and Conway, 2010, p. 186). This was hardly the only example of counter-lobbying by industry funded think tanks. Indeed, Jacques et al. (2008) study 141 English-language books that adopt a sceptical stance on environmental issues and find that 92% of them are associated with conservative think tanks. Dunlap and Jacques (2013) study books denying climate change in particular, and find that at least 90% of them do not undergo peer review. One of the key funders of climate denial efforts was ExxonMobil, which has given over \$20 million to sow doubt about global warming to a wide array of think tanks, with the biggest recipient being the Competitive Enterprise Institute with \$2 million of ExxonMobil support (Hoggan and Littlemore, 2009, p. 82). After its support for climate denialists was exposed publicly by Greenpeace, ExxonMobil cut off funding to the Competitive Enterprise Institute, though it nevertheless continued to donate over \$2 million to other denialists in 2007 (Hoggan and Littlemore, 2009, p. 84).

¹Since the founding of the Atlas Economic Research Foundation in 1981, the number of free market think tanks has grown rapidly and now tops 400 (Djelic, 2014).

There is a rich literature in economics and political science on the strategies used by interest groups to influence the political process. The literature has focused on two main channels of influence: campaign contributions (Austen-Smith, 1995, 1998; Bernheim and Whinston, 1986; Snyder, 1990; Groseclose and Snyder, 1996; Kroszner and Stratmann, 1998; Cotton, 2009) and informational lobbying (Crawford and Sobel, 1982; Potters and van Winden, 1992; Lohmann, 1995; Krishna and Morgan, 2001). Grossman and Helpman (2001) provide an excellent overview of both strategies. Recent models have begun to incorporate both options. For example, Bennedsen and Feldmann (2006) study the choice between lobbying and campaign contributions; they assume the politician can see whether an interest group invested in gathering information, so a failure to lobby is informative, and hence there is a tradeoff between lobbying and campaign contributions, with interest groups specializing in one or the other. Dahm and Porteiro (2008) have a similar model to Bennedsen and Feldmann, but they assume the politician cannot observe whether an interest group has engaged in information gathering or not, and hence they find that interest groups invest in both strategies.

In this paper, we introduce a new channel through which an interest group can influence policy: supplying information about the credibility of political rivals. We develop a model that extends the literature on informational lobbying (Grossman and Helpman, 2001) to explore the impact of doubt-creation strategies on the public policy process. We identify conditions under which the creation of doubt can weaken regulatory policy or block its passage altogether. Because a firm that faces increased costs from regulation is typically not a credible source of information, it may need to resort to indirect tactics to influence public policy. In particular, we model the two tactics described above, motivated by recent literature on the role of business in the policy process. First, we model the process of undermining the credibility of the NGO as a form of Bayesian persuasion, not about the state of the world itself but about the political bias of the NGO. Second, we model the operations of a “think tank” that can provide noisy information about the state of the world that is contrary to that provided by the NGO. We show that the two mechanisms can be used to produce identical effects, and we characterize the conditions under which they are effective.²

²One might think that the NGO should be able to engage in Bayesian persuasion on its own behalf, to counter the actions of the firm. However, models of Bayesian persuasion are symmetric information models, while the NGO in our model has

We make two main contributions.³ First, we show how a public signal (Bayesian persuasion) can be used not just to persuade decision makers about the state of the world, as in Kamenica and Gentzkow (2011), but also to persuade decision makers about the *credibility* of other players who are offering information. This type of “meta-communication” is becoming increasingly important as information overload leads people to seek out trusted sources to filter information for them. Naturally, such persuasion is profitable for the firm, but we show that when making policy is costly the policy maker also benefits from persuasion. However, the use of Bayesian persuasion requires that a firm be able to commit to making the results of its “investigation” into NGO bias public, even when those results are unfavorable to its interests. When this is not possible, the firm may seek an alternative mechanism through which to obtain the Bayesian persuasion benchmark.

Second, we show that an industry-oriented think tank can implement the Bayesian persuasion benchmark, by acting as a possible radical instead of a credible moderate. That is, the think tank lobbies the policymaker in order to communicate its view of the state of the world. Instead of developing a reputation for truthful communication, however, the think tank intentionally obscures its political bias and acts as a radical some of the time. The mechanism for doing so involves carefully selecting a mix of employees that vary along the spectrum from moderate to extreme. To the extent that firing fellows is costly or time-consuming, the think tank can serve as a commitment mechanism for doubt creation efforts. We show how the design of a think tank’s political bias optimally responds to the characteristics of the issues it addresses and the NGO private information about its own type. Thus, the conditions for the use of Bayesian persuasion do not obtain. A similar situation would arise in the example of a prosecutor persuading a judge in Kamenica and Gentzkow (2011). The prosecutor and the judge have symmetric information—neither knows for sure whether the defendant is guilty or not. But the defendant knows whether he is guilty or not, and his asymmetric information violates the conditions for the use of Bayesian persuasion. Moreover, in our model there is no point for the NGO to engage in Bayesian persuasion regarding the type of the firm, as it will be common knowledge that the firm has an extreme form of bias, i.e. it always prefers no policy be taken regardless of the state of the world.

³We also contribute to the literature on strategic communication of “soft” or unverifiable information, extending the canonical Grossman-Helpman (2001) model of informational lobbying to incorporate the realistic assumption that policymakers face a cost of formulating and implementing policy. We show that this change makes the policymaker more skeptical in the sense that NGOs must be more closely aligned with the policymaker’s interests if they wish to credibly convey information.

opposition it faces.

Two other recent papers address related phenomena from different perspectives. Stone (2011) develops a model in which a special interest group covertly funds research that it releases into the public domain, with the hope that the research will appear unbiased to a policymaker who happens upon it randomly. This strategy aims to crowd out other, unsponsored, research and relies upon the covert nature of its funding for success. Shapiro (2016) develops a model in which a journalist reports the state of the world to a voter who chooses a policy. The journalist's desire to maintain a reputation for being unbiased can lead him to adopt a policy of "balanced reporting," in which he reports claims by special interests even when he knows they are not true. The mechanisms we study are very different from, though complementary to, the ones in these papers. Unlike Stone (2011), the interest group in our model does not rely on covert funding of research; instead it commits to a fully public investigation. Unlike Shapiro (2016), the interest group in our model need not communicate through a journalist, so issues of journalistic reputation maintenance do not arise; instead, the interest group communicates directly to a policymaker. Unlike either paper, the interest group in our model can achieve its goals by directly undermining the credibility of a science-based NGO, rather than offering its own scientific evidence.

This article is also part of a growing body of work on corporate nonmarket strategy in response to regulatory threats. Early work in this vein showed that firms may engage in substantive self-regulation in order to preempt regulatory threats, and that under full information this can be socially beneficial (Segerson and Miceli, 1998; Maxwell, Lyon and Hackett, 2000; Fleckinger and Glachant, 2011).⁴ More recently, researchers have pointed out that firms may prefer to avoid taking costly substantive action, and instead use persuasive mechanisms to deflect regulatory pressures and influence public debate.⁵ One approach is to influence the underlying preferences of the public through indirect lobbying (Yu, 2005; Prieur and Zou, 2017). An alternative approach is to influence the beliefs of stakeholders, without changing their

⁴Related research studies the possibility that NGOs may campaign against a firm even after the regulatory process has been settled in the firm's favor (Daubanes and Rochet, 2016).

⁵Millner and Ollivier (2016) argue that the processes by which public opinion is formed deserve much more attention from economists.

preferences. One such technique is for a firm to disclose a large amount of low-quality information that is difficult for stakeholders to evaluate (Sinclair-Desgagné and Gozlan, 2003). A variant of this technique is to spend money to generate and publicize biased scientific findings, in the presence of scientific uncertainty, in order to take advantage of citizens’ misperceptions (Bramoullé and Orset, 2018). A second technique is for firms to covertly fund front groups, also known as “astroturf lobbying” groups, which pose as authentic grassroots networks of citizens in order to influence public policy (Lyon and Maxwell, 2004; Cho et al., 2011; Walker, 2014). A third technique is greenwashing, which commonly involves the selective disclosure of positive information about a firm’s environmental performance while neglecting to mention negative information (Lyon and Maxwell, 2011; Lyon and Montgomery, 2013).⁶ Each of the foregoing techniques faces limitations on the conditions under which it is effective: disclosing low-quality information may fail with skeptical stakeholders, astroturf lobbying cannot silence an existing NGO with opposing preferences, and greenwashing may be too risky for firms with poor track records of corporate citizenship. We show that doubt creation is attractive because it can be used by any type of firm regardless of its track record. However, it is only effective when the costs of implementing policy fall within a certain range. Hence, firms need access to a suite of possible information-based nonmarket strategies that respond to specific contexts.

The remainder of the article is structured as follows. Section 2 presents the basic model of lobbying by an NGO with known bias when policy making is costly. Section 3 extends this to the case of an NGO whose bias is unknown to other players. Section 4 introduces the firm and its role in the creation of doubt about the NGO’s credibility. Section 5 studies the creation of a “think tank” by the firm in order to oppose the NGO’s position. Section 6 offers conclusions. Proofs are presented in the appendix.

2 A Model of Lobbying with Costly Policymaking

In this section, we present the basic model of informational lobbying when policymaking is costly. This model is identical to that of Grossman and Helpman (2001, chapter 5.1) except that we make the realistic

⁶As summarized by Lyon and Montgomery (2015), there is a large literature in management on the deceptive practices firms use to fend off pressures for change from social movements and public policy.

assumption that the policymaker has a fixed cost F of formulating and implementing policy, which is common knowledge.⁷ We show that a key implication of costly policymaking is to make the policymaker more skeptical, in the sense that NGOs must be more closely aligned with the policymaker's interests if they wish to credibly convey information.

The notion that it is costly to introduce a policy is an intuitive one. Legislators are severely time-constrained and under-staffed (Drutman, 2015). In the United States, legislators spend as much as half of their time simply raising funds for the next electoral campaign (Roemer, 2015). Indeed, it is common for legislators to vote on complex bills without even having read through their contents, because they simply do not have time to read everything they vote on. Thus, devoting the time and effort to introduce a piece of legislation carries with it substantial opportunity costs. Regulatory actions carry with them perhaps even greater costs. In the United States, legislative procedures require that regulators follow a detailed set of steps before issuing a regulation, including putting forward a Notice of Proposed Rulemaking (NPR), allowing interested parties to comment on the NPR, and making detailed responses to all substantive comments received (McCubbins, Noll and Weingast, 1987). In addition, Executive Order 12866 requires that if a regulation is expected to impose substantial costs on industry (\$100 million or more), then regulators must conduct a Regulatory Impact Analysis that demonstrates to the Office of Management and Budget that the benefits of the rule exceed the costs.⁸ Thus, the costs of imposing any new regulation are quite substantial, both in terms of time and in terms of real agency resources (Viscusi, Vernon and Harrington, 1998, pp. 321-322.) We follow the literature in assuming that there is a fixed cost of proposing new policies (Glazer and McMillan 1990; Glazer and McMillan 1992; Maxwell, Lyon and Hackett 2000).⁹ To the best of our knowledge, however, prior work using the interest-group signaling model (e.g., Grossman and Helpman, 2001) has ignored this aspect of the political arena.

⁷In future work it may be an interesting extension to allow for uncertainty about F .

⁸For details, see https://www.whitehouse.gov/sites/default/files/omb/assets/OMB/circulars/a004/a-4_FAQ.pdf.

⁹One can imagine other formulations in which policymaking has both fixed and variable costs, but we leave such an extension for future research.

2.1 Model Elements

There is a social issue whose severity is as yet unknown, and may turn out to have either a high impact or a low impact on society. (For concreteness, we henceforth refer to the issue as an “environmental” issue, but the case studies in Oreskes and Conway (2010) clearly show that the same mechanisms arise for health or other social issues.) The state of the world is θ_i $i \in \{L, H\}$, a real-valued positive number that indicates whether the needed policy stringency is low or high. The state of the world takes on the value θ_L with probability $1 - q$ and $\theta_H > \theta_L$ with probability q ; the expected state of the world is then $\bar{\theta} = q\theta_H + (1 - q)\theta_L$. Legislation of stringency p , also real-valued and positive, may be passed to ameliorate the losses from the environmental harm. There are three actors: a policy maker (PM, whom we indicate with a subscript M) and an environmental NGO (whom we will indicate with a subscript E)—each of whom seeks to minimize the difference between the policy and their own ideal point—and a firm, which seeks to minimize policy stringency.

The PM’s problem is to

$$\max_p U_M(p, \theta) = -(p - \theta)^2 - F.$$

Thus, if it acts, it prefers to set $p = \theta$. However, because the PM also has a fixed cost F of undertaking the effort to create a policy, it may choose not to act in some states of the world. This will turn out to have important effects on the firm’s strategy in later sections. The cost of policymaking is an agency cost that creates a wedge between the PM’s incentives and any aggregate measure of social welfare. Thus, our analysis should be viewed as a positive rather than a normative one.

The NGO’s net utility is

$$U_E(p, \theta, \delta) = -(p - \theta - \delta)^2 - L,$$

where δ is a measure of the NGO’s ideological position or bias, relative to that of the PM, and L is the exogenous lobbying cost incurred by the NGO. The NGO’s willingness to expend resources thus may serve as a signal of the severity of the state of the world. Its preferred policy is equal to $\theta + \delta$. The NGO may be either a “moderate” with a low bias or a “radical” with a high bias (the threshold that divides the two is

defined by the characteristics of the problem, as described below). Regardless, whatever its bias, the NGO always prefers a more stringent policy than does the PM, that is, $\delta > 0$.

The firm's objective function is different from that of the other players: it simply wants to minimize the policy burden that it faces, as in Lyon and Maxwell (2004). The firm in our model is not driven by altruistic motives, nor by the personal ideological preferences of managers; it is strictly profit-maximizing. Its most preferred situation is when the policymaker takes no action at all. As a result, the firm is not a credible source of information, since it always has incentives to claim that the state is low. This lack of credibility is precisely what motivates the firm to utilize indirect methods of influence. The firm's payoff function is simply

$$\Pi(p) = -cp^2.$$

This is a simple and natural formulation in which the cost to the firm of complying with policy is increasing and convex in policy stringency, as is commonly assumed in models of regulatory compliance (Fleckinger and Glachant, 2011; Heyes and Martin, 2017).

The basic game involves three stages:

1. The NGO learns the state of the world θ .
2. The NGO can lobby, *i.e.* expend resources to signal its information to the PM, and the PM updates its beliefs.
3. The PM decides whether to incur the fixed cost F of going through a policymaking process, and if so it makes a policy decision p and the game ends.

The game is solved via backwards induction to ensure a subgame perfect Nash equilibrium.

2.2 Policy Decisions

Suppose initially that the PM has no information about the state of the world. If the PM implements no policy at all, its expected utility is

$$EU_M^{NP} = -(1-q)(\theta_L)^2 - q(\theta_H)^2,$$

where E indicates the expectations operator. If the PM takes action without information about the state it maximizes its expected utility by setting

$$p = \bar{\theta}$$

and obtains expected utility

$$EU_M^{UI} = -(1-q)q(\theta_H - \theta_L)^2 - F.$$

Thus it is worthwhile for the PM to take action if $EU_M^{UI} > EU_M^{NP}$ or

$$F < F_q \equiv \bar{\theta}^2. \tag{1}$$

If the PM learns the underlying state of the world θ before choosing a policy, then if it acts it will choose $p_i = \theta_i$, for $i \in \{L, H\}$. Its utility in either state is then $U_M = -F$. If it opts not to act in state i , its utility is $-(\theta_i)^2$. Thus, the informed PM acts in both states if $F < F_L \equiv (\theta_L)^2$, never acts if $F > F_H \equiv (\theta_H)^2$, and acts if and only if the state is high if $F \in (F_L, F_H)$.

2.3 Lobbying by the NGO

The NGO has a bias of $\delta > 0$ that is common knowledge. The NGO must weigh the costs and benefits of lobbying to inform the PM. (If there were no cost of lobbying, then the NGO would always engage in uninformative “cheap talk.”) As in the standard model of Grossman and Helpman (2001, chapter 5), because lobbying is costly there is the possibility of an informative lobbying equilibrium in which the NGO appears before the PM in one state and stays home in the other state.¹⁰ We focus on pure-strategy lobbying

¹⁰As is also true there, the equilibrium is not unique. In particular there is always a “babbling equilibrium” in which the PM always sets $p = q\theta_H + (1-q)\theta_L$ and the NGO has no incentive to lobby truthfully. We follow the literature in focusing on the informative equilibrium.

equilibria.

To characterize the range of biases for which informative lobbying is possible, define $\underline{\delta} \equiv L/2\theta_H + \theta_H/2 - \theta_H$ and $\bar{\delta} \equiv L/2\theta_H + \theta_H/2 - \theta_L$, as well as $\underline{\delta} \equiv L/2(\theta_H - \theta_L) - (\theta_H - \theta_L)/2$ and $\bar{\delta} \equiv L/2(\theta_H - \theta_L) + (\theta_H - \theta_L)/2$.

We can then characterize the impact of the PM's fixed cost on the lobbying equilibrium in the following proposition.

Proposition 1 *When $F < F_L$, an informative lobbying equilibrium exists for $\delta \in (\underline{\delta}, \bar{\delta})$. When $F \in [F_L, F_H)$, an informative lobbying equilibrium exists for $\delta \in (\underline{\delta}, \bar{\delta})$. The range of NGO biases that support an informative equilibrium is of the same magnitude in both cases, i.e. $\bar{\delta} - \underline{\delta} = \bar{\delta} - \underline{\delta}$, but $\bar{\delta} < \bar{\delta}$.*

Thus, when policymaking is costly ($F > F_L$) the PM becomes more skeptical than when $F = 0$ in the sense that NGOs must be more closely aligned with the PM's objectives in order to be deemed credible. Having established this key implication of costly policymaking, throughout the remainder of the paper we focus on the case where $F \in (F_L, F_H)$. Thus, we avoid both the case where the PM always acts (we show later in the paper that Bayesian persuasion is never profitable for the firm in this case) and the case where the PM never acts (persuasion is not needed in this case), and focuses attention on the interesting case in which the PM will not act if it believes the state is low.

3 Costly Policy Making when the NGO's Bias is Unknown

Suppose now the PM is unsure whether the NGO is entirely reliable, that is, whether it might be a “radical” with a high bias rather than a “moderate” with a smaller bias. For simplicity, we model this by assuming that δ takes one of two values. With probability ζ , the NGO's bias takes on the value $\delta = R > \bar{\delta}$, which is so large that the NGO always has incentives to lobby regardless of the state (the NGO is a “radical”). With probability $1 - \zeta$, the NGO's bias takes on the value $\delta = M \in (\underline{\delta}, \bar{\delta})$, so that the NGO lobbies if and only if the state is θ_H (the NGO is a “moderate”).¹¹ Facing this uncertainty, the PM may not be able to

¹¹Another modeling option would be to assume that the bias of the NGO is unknown but that its distribution is common knowledge, e.g. δ is distributed uniformly on $(\underline{\delta}, \Delta)$ with $\Delta > \bar{\delta}$. However, this approach is more cumbersome, so we have opted

rely upon the NGO as a means of learning the underlying state of the world θ before choosing a policy.

3.1 The Policymaker's Inferences and Actions

Conditional upon observing the NGO lobby, the PM updates its belief that the state is actually high. The unconditional probability that the NGO lobbies is now

$$\Pr(\text{NGO lobbies}) = q + (1 - q)\zeta.$$

Using Bayes' Rule, the PM's updated belief ρ conditional on observing that the NGO lobbies is

$$\rho = \Pr(\theta = \theta_H | \text{NGO lobbies}) = \frac{q}{q + (1 - q)\zeta}. \quad (2)$$

Note that $\rho > q$ because $\zeta \in (0, 1)$, so NGO lobbying still leads the PM to increase its belief that the state is high, although it does not believe the state is high with certainty. With this updated belief, if the PM implements a policy conditional on observing that the NGO lobbies, it will set

$$p_\rho = \rho p_H + (1 - \rho)p_L = \rho\theta_H + (1 - \rho)\theta_L. \quad (3)$$

If the NGO does not lobby, the PM can still infer with certainty that the state is L , so it sets

$$p_L = \theta_L.$$

Define

$$\delta_U \equiv \frac{L}{2p_\rho} + \frac{p_\rho}{2} - \theta_H$$

and

$$\delta^U \equiv \frac{L}{2p_\rho} + \frac{p_\rho}{2} - \theta_L.$$

Then the logic of the proof of Proposition 1 can be used to show that an (imperfectly) informative equilibrium exists for the appropriate range of NGO biases.

Lemma 2 *When the NGO's bias is uncertain, an imperfectly informative lobbying equilibrium exists for $\delta \in (\delta_U, \delta^U)$. In this equilibrium, the radical NGO lobbies in both states of the world but the moderate NGO*

for the simpler one presented in the text.

lobbies only when the state is high. The policymaker sets p_ρ when the NGO lobbies and sets p_L when the NGO does not lobby.

By definition, then, a moderate NGO must have a bias $\delta = M$ in the range (δ_U, δ^U) . It is easy to show that as $\zeta \rightarrow 0$ the above condition converges to $\delta \in (\underline{\delta}, \bar{\delta})$, the condition for informative lobbying when the NGO's bias is known. Thus, as long as ζ is small the NGO will still lobby informatively.

Because $F > F_L$, the PM does not act if the NGO does not lobby. Even if the NGO does lobby, however, uncertainty about the NGO's bias alters the PM's calculus regarding whether to take action, as shown in the next lemma.

Lemma 3 *If the policy maker is uncertain of the NGO's bias, it takes action if and only if*

$$\zeta < \zeta^{NP} \equiv \frac{q}{1-q} \frac{(\theta_H - \sqrt{F})}{(\sqrt{F} - \theta_L)}, \quad (4)$$

Note that $\zeta^{NP} < 1$ is equivalent to $F > \bar{\theta}^2$. We can also express ζ^{NP} in terms of the PM's posterior belief ρ regarding the state of the world, so that the PM only takes action if

$$\rho > \rho^{NP} = \frac{\sqrt{F} - \theta_L}{(\theta_H - \theta_L)}, \quad (5)$$

or in terms of F , with the PM only taking action if

$$F < F_\rho = (p_\rho)^2. \quad (6)$$

The parameter value ζ^{NP} will be critical in the firm's choice of doubt-creation strategy in section 4. As intuition would suggest, if the PM is sufficiently skeptical of the NGO's credibility, then the PM will not take action even if the NGO lobbies. Moreover, if $F > F_\rho$ then the PM will never act, even when the NGO lobbies. Thus in what follows we focus on the case where $F < F_\rho$.

3.2 The Effect on Profits of Uncertainty about the NGO's Bias

Does the firm benefit when the PM is uncertain about the NGO's type, i.e. when it faces a "noisy" NGO? It might at first seem obvious that the answer would be "yes," but in fact the firm faces a tradeoff when

the NGO is noisy, even in the simple case where $F < F_L$. On one hand, the firm benefits from noise because it never obtains the worst policy, p_H , and instead obtains $p_\rho < p_H$ when the NGO lobbies. On the other hand, it is less likely to obtain the least stringent policy p_L , which now occurs with probability $(1-q)(1-\zeta) < (1-q)$. Thus, it is not obvious whether the firm gains from uncertainty about the NGO's bias.

When $F > F_L$, the PM never implements p_L , and expected profits as a function of ζ are

$$E\Pi^{IP} = -\frac{c}{q+(1-q)\zeta}(qp_H+(1-q)\zeta p_L)^2,$$

with $E\Pi^{IP}|_{\zeta=0} = -cqp_H^2$ and $E\Pi^{IP}|_{\zeta=1} = -c\bar{p}^2$. Substituting in the PM's optimal policies $p_H = \theta_H$ and $p_L = \theta_L$ and differentiating with respect to ζ yields

$$\frac{\partial(E\Pi^{IP})}{\partial\zeta} = c\frac{(1-q)q^2}{(q+(1-q)\zeta)^2}(\theta_H-\theta_L)^2 - c(1-q)\theta_L^2.$$

A bit of algebra shows that expected profits are strictly decreasing in ζ if $\theta_H < 2\theta_L$, are strictly increasing in ζ if $\theta_H > \theta_L(1+q)/q$, and have an interior maximum in ζ if $\theta_H \in (2\theta_L, \theta_L(1+q)/q)$. Expected profits are strictly concave in ζ because

$$\frac{\partial^2(E\Pi^{IP})}{\partial\zeta^2} = -\frac{2c(1-q)^2q^2(\theta_H-\theta_L)^2}{(q+(1-q)\zeta)^3} < 0. \tag{7}$$

Regardless of the precise shape of the expected profit function when $F > F_L$, there is a discontinuity in expected profits at $\zeta = \zeta^{NP}$, as implied by Lemma 3. The reason is that for $\zeta \geq \zeta^{NP}$ the PM takes no action even if the NGO lobbies, and hence expected profits become $E\Pi^{IP} = 0$. Thus, the firm always prefers to have $\zeta \geq \zeta^{NP}$, but whether it benefits from a marginal increase in ζ when $\zeta < \zeta^{NP}$ depends upon parameter values. The shape of the expected profit function for the case with $\theta_H < 2\theta_L$ is shown in bold in Figure 1.

4 Strategic Creation of Doubt

We turn now to the possibility of strategic creation of doubt by the firm. How can the firm create doubt about the NGO's bias? One modeling option would be to assume the firm (or an agent hired by the firm)

BlackandWhite




Figure 1: Expected Profits in Case (a)

expends resources to search for hard, verifiable, evidence that the NGO is an extremist. Such evidence could only be found if the NGO truly were a radical, and so would provide simple, credible evidence to the policymaker that the NGO should be ignored. There would be no need for the policymaker to engage in complicated inferences regarding the NGO, and the firm could wait until it observes whether the NGO lobbies before releasing its evidence. The problem with this modeling approach, however, is that it simply does not explain the practices described in Hoggan and Littlemore (2009) or Oreskes and Conway (2010). Those books detail efforts to discredit NGOs who were moderate scientists, and not radical extremists. They describe attacks made with soft, unverifiable, claims, such as that Ben Santer had doctored the IPCC report “to deceive policy makers and the public.” This was an assertion about Santer’s motives, not a verifiable piece of hard information.

Because the attacks on NGOs in practice are typically not verifiable (and usually not true), we adopt a modeling approach that involves the provision of soft, unverifiable, information about the NGO. Given the

firm’s payoff function, it cannot credibly engage in cheap talk simply asserting that the NGO is a radical—the firm has an incentive to claim the NGO is radical regardless of whether it is true. If Exxon attacks scientists for being radical, few policymakers or members of the public will grant this much credence. And as mentioned above, the firm typically cannot obtain verifiable information that an NGO is an extremist. If the firm can produce some minimal amount of evidence, however, it has a much better chance of being believed.

Suppose the firm can fund an “investigation” into the bias of the NGO, one that need not necessarily be neutral or objective, but one that the firm commits to make public. The PM’s initial belief is that the NGO is a radical with probability ζ^0 . As long as the investigation meets the minimal requirement of Bayes plausibility, namely that the expected posterior distribution over the NGO’s bias is equal to the prior distribution ζ^0 , then the firm may be able to engage in Bayesian persuasion (Kamenica and Gentzkow, 2011).¹² Formally, Bayesian persuasion implies that the firm’s investigation creates a signal π which consists of a realization space Y , and a family of conditional probability distributions $\pi(\cdot|\delta)$ over Y . The firm can choose the probability distributions over the realization in any way that is Bayes plausible. For simplicity, let the signal realization be binary, so that either $y = M$ (for moderate) or $y = R$ (for radical).

Our model differs from Kamenica and Gentzkow (2011) in that we have two very different sources of information transmission. The first is the signal π sent by the firm (which conveys information regarding the NGO’s credibility), and the second is an action taken by the NGO (which conveys information regarding the state of the world), whereas Kamenica and Gentzkow (2011) deal only with a single signal about the

¹²Admittedly, the assumption that the firm can commit to make the results of its investigation public regardless of their outcome is a strong one, since *ex post* the firm might prefer to reveal only “findings” suggesting that the expert is an extremist. However, one can envision settings where such commitment is possible; for example, the firm might contract with a third party to hack into a private email server and reveal all the illegally acquired emails publicly via Wikileaks. Moreover, Kamenica and Gentzkow (2010, pp. 2598-2599) argue that the commitment assumption can be dispensed with if the Sender of information sends verifiable signals that are supersets of the Sender’s own privately observed signals. At any rate, this same critique can be leveled against many applications of Bayesian persuasion. One can also think of the Bayesian persuasion model as a benchmark of what can be accomplished by someone attempting to cast doubt on the credibility of an expert. In section 5 we show how the firm can accomplish the same goals using an alternative approach whose commitment value may be viewed as more plausible.

state of the world. It is important to recognize that the two forms of information transmission in our model have very different structural forms. In effect, we have two different games going on within the model, each with very different informational properties. The game of Bayesian persuasion is a game of symmetric but incomplete information between the firm and the PM. In this game, the firm can costlessly commit to transmit a family of conditional probability distributions with public realizations that convey information about the bias of the NGO. The lobbying game is a game of asymmetric information between the NGO and the PM. In this game, the NGO cannot commit to a strategy, but can take a costly action (lobbying) that conveys information about the state of the world.

Let $\widehat{\zeta}_y(\delta = R|y)$ be the posterior belief of the PM that the NGO is a radical, conditional on observing the realization y of the firm's signal. For notational simplicity, we will write $\widehat{\zeta}_y = \widehat{\zeta}_y(\delta = R|y)$. Then $\pi(y|\delta)$ is the conditional probability of observing the realization y when the true state is δ . For notational simplicity, we will write $\pi_\delta = \pi(R|\delta)$. Let λ be the unconditional probability of observing the realization $y = R$ and $1 - \lambda$ be the unconditional probability that $y = M$. Then

$$\lambda = \pi(R|\delta = R)\zeta^0 + \pi(R|\delta = M)(1 - \zeta^0) = \pi_R \zeta^0 + \pi_M(1 - \zeta^0). \quad (8)$$

Bayesian updating requires that for signal realization $y = M$, the PM's posterior belief about the NGO's bias is

$$\widehat{\zeta}_M = \frac{(1 - \pi_R) \zeta^0}{(1 - \pi_R) \zeta^0 + (1 - \pi_M)(1 - \zeta^0)}, \quad (9)$$

and for realization $y = R$ it is

$$\widehat{\zeta}_R = \frac{\pi_R \zeta^0}{\pi_R \zeta^0 + \pi_M(1 - \zeta^0)}. \quad (10)$$

Bayes plausibility requires that the expected posterior distribution over the NGO's bias is equal to the prior distribution, that is,

$$\widehat{\zeta}_R \lambda + \widehat{\zeta}_M(1 - \lambda) = \zeta^0. \quad (11)$$

The modified game now involves four stages:

1. The firm designs an experiment and sends a signal about the bias of the NGO, and both the signal and its realization are observed by the PM.

2. The NGO learns the state of the world θ .
3. The NGO can lobby, *i.e.* expend resources to convey its information to the PM, and the PM updates its beliefs.
4. The PM decides whether to incur the fixed cost F of going through a policymaking process, and if so it makes a policy decision p and the game ends.

In the modified game, the PM observes multiple pieces of information. From the firm, it observes both the firm’s choice of signal (the design of its experiment), as well as the signal realization regarding the NGO’s bias. From the NGO, the PM observes the NGO’s choice of lobbying action, which conveys information about the state of the world.

4.1 An Example

A simple example may help to clarify the role of Bayesian persuasion in our setting. Suppose the PM has a prior belief that there is a 10% chance the NGO is a radical, so $\zeta^0 = 0.1$. The firm structures its investigation by choosing π_R and π_M , knowing that the PM will update its beliefs according to Bayes’ Rule. Perhaps the firm hires a hacker to break into the emails of a prominent climate scientist and then interprets the results in the most negative possible light.¹³ Note that Bayes plausibility means it is pointless for the firm to choose an investigation that will always says the NGO is a radical, no matter what. Such an investigation would imply that $\pi_R = \pi_M = 1$, that $\lambda = 1$, and that $\hat{\zeta}_M = 0$ and $\hat{\zeta}_R = \zeta^0$. Hence it would have no impact on the PM’s beliefs, as the PM will always observe the signal realization indicating the NGO is a radical and will not update its beliefs. However, suppose the firm “rigs” its investigation so that if the NGO really is a moderate, the investigation says “moderate” with 50% probability and “radical” with 50% probability (that is, $\pi_M = 0.5$), while if the NGO really is a radical the investigation always says “radical” (that is, $\pi_R = 1.0$). This means the investigation says the NGO is a radical with probability

¹³This was essentially the case of the so-called “Climategate” scandal, although it has never been established who directed the hack. https://en.wikipedia.org/wiki/Climatic_Research_Unit_email_controversy#Public_opinion_and_political_fallout

$\lambda = 0.1 * 1.0 + 0.9 * 0.5 = .55$, even though there is only a 10% chance the NGO really is a radical. The PM sees that the investigation is rigged, so if it observes the signal realization that the NGO is a radical, it only updates its belief to $\hat{\zeta}_R = (1.0 * 0.1) / .55 = 0.1818$. If the PM observes the signal realization that the NGO is a moderate, it updates its belief to $\hat{\zeta}_M = 0 / .45 = 0$. In other words, if the signal ever says the NGO is a moderate, the PM can be sure this is true. Finally, Bayes plausibility is satisfied because $.1818 * .55 + 0 * .45 = 0.1$. Thus, the firm is free to concoct a highly biased investigation, and the PM pays attention to the signal produced, but the PM is not fooled by the investigation. Nevertheless, the biased investigation may be useful to the firm if the PM makes a policy choice that is more favorable to the firm when the signal realization says the NGO is a radical.

4.2 Is Bayesian Persuasion Profitable?

To determine whether Bayesian persuasion is valuable, it is sufficient to construct the concave closure of $E\Pi^{IP}(\zeta)$, which we will denote by $C(E\Pi^{IP}(\zeta))$ and to assess whether $C(E\Pi^{IP}(\zeta^0)) > E\Pi^{IP}(\zeta^0)$.¹⁴ Note that inequality (7) shows that $E\Pi^{IP}(\zeta)$ is strictly concave, which suggests that $C(E\Pi^{IP}(\zeta)) = E\Pi^{IP}(\zeta)$ for all ζ , and that there would be no value to Bayesian persuasion. However, when $F > F_L$ and $\zeta^{NP} < 1$ (a necessary and sufficient condition for both of these to hold is $F > \bar{\theta}^2$), there is an upward jump discontinuity in expected profits at ζ^{NP} . In this case, it is possible to construct a concave closure of the $E\Pi^{IP}$ function with $C(E\Pi^{IP}(\zeta)) > E\Pi^{IP}(\zeta)$ for some values of ζ^0 .

The construction of the concave closure naturally depends upon the shape of $E\Pi^{IP}$, which section 3.2 showed can be strictly decreasing in ζ , strictly increasing in ζ , or have an interior maximum. Begin with the first possibility ($\theta_H < 2\theta_L$), which we refer to as case (a) and which is illustrated in Figure 2, since it is the simplest. The dashed bold line forms the concave closure of $E\Pi^{IP}(\zeta)$, and it is easy to see that the concave closure connects $(0, -cq(p_H)^2)$ and $(\zeta^{NP}, 0)$. Because $E\Pi^{IP}$ is strictly concave for $\zeta < \zeta^{NP}$, $C(E\Pi^{IP}(\zeta)) > E\Pi^{IP}(\zeta)$ for all $\zeta < \zeta^{NP}$ and Bayesian persuasion is profitable for $\zeta^0 \in (0, \zeta^{NP})$.

It is also easy to see from Figure 2 that the concave closure will retain this shape even if the expected

¹⁴See Corollary 2 of Kamenica and Gentzkow (2011).




Figure 2: Concave Closure of Expected Profits for Case (a)

profit function is upward-sloping at $\zeta = 0$, as long as the slope of the expected profit function is less than the slope of the line segment connecting $(0, -cq(p_H)^2)$ and $(\zeta^{NP}, 0)$.¹⁵ Formally, when $\theta_H > 2\theta_L$, the shape of the concave closure remains as shown in Figure 2 as long as

$$\left. \frac{\partial (E\Pi^{IP})}{\partial \zeta} \right|_{\zeta=0} \leq \frac{cq(p_H)^2}{\zeta^{NP}},$$

which some algebra shows is equivalent to $\sqrt{F} > \theta_H/2$ or equivalently $F > F_H/4$. We refer to this as case (b).

Alternatively, suppose $\theta_H > 2\theta_L$ and

$$\left. \frac{\partial (E\Pi^{IP})}{\partial \zeta} \right|_{\zeta=0} > \frac{cq(p_H)^2}{\zeta^{NP}},$$

or equivalently, $F < F_H/4$, which defines case (c). In this case, the line segment constructed in the first case cannot form the concave closure of $E\Pi^{IP}$ because it lies beneath the $E\Pi^{IP}$ function for at least a portion of the range $(0, \zeta^{NP})$. In this case, in order to construct the relevant concave closure it is necessary to find

¹⁵Note that the three cases we identified in section 3.2 regarding the shape of the expected profit function reduce to two cases here.

BlackandWhite




Figure 3: Concave Closure of Expected Profits in Case (c)

a line segment that is tangent to $E\Pi^{IP}$ and passes through $(\zeta^{NP}, 0)$. (Such a tangency exists even in the case where $E\Pi^{IP}$ is strictly increasing in ζ because $E\Pi^{IP} < 0$ for all ζ .) This case is illustrated in Figure 3. The concave closure of expected profits is shown by the dashed bold line, which is constructed so as to be anchored at the right-hand end at $(\zeta^{NP}, 0)$ and to be tangent to $E\Pi^{IP}$ at the left-hand end. From the graph it is clear that persuasion creates value as long as $\zeta^0 \in (\zeta^T, \zeta^{NP})$.

The tangency occurs at the value ζ^T that solves

$$\frac{0 - E\Pi^{IP}(\zeta)}{\zeta^{NP} - \zeta} = \frac{\partial (E\Pi^{IP})}{\partial \zeta}. \quad (12)$$

Substituting $E\Pi^{IP}$ and $\partial (E\Pi^{IP}) / \partial \zeta$ into (12) and simplifying yields

$$(1-q)^2\theta_L[2q\theta_H + (1-q)\theta_L\zeta^{NP} - q\theta_L]\zeta^2 + 2q(1-q)[(1-q)\theta_L^2\zeta^{NP} + q\theta_H^2]\zeta + q^2\theta_H[q\theta_H - \zeta^{NP}(1-q)(\theta_H - 2\theta_L)] = 0.$$

Applying the quadratic formula reveals that there is one positive root, which is

$$\zeta^T = \frac{q(\theta_H - 2\sqrt{F})}{(1-q)(2\sqrt{F} - \theta_L)}. \quad (13)$$

Clearly ζ^T is a function of q, θ_L, θ_H , and F but we will suppress that dependence to reduce notational clutter. Thus we have demonstrated the following proposition.

Proposition 4 *Bayesian persuasion is profitable if $\zeta^{NP} < 1$ (equivalently, $F > \bar{\theta}^2$) and (a) $\theta_H < 2\theta_L$ and $\zeta^0 < \zeta^{NP}$, (b) $\theta_H > 2\theta_L$ and $F > F_H/4$ and $\zeta^0 < \zeta^{NP}$, or (c) $\theta_H > 2\theta_L$ and $F < F_H/4$ and $\zeta^0 \in (\zeta^T, \zeta^{NP})$.*

It is clear from Figures 2 and 3 that the fixed cost of policymaking is essential to making Bayesian persuasion valuable in a lobbying model such as this. The role of persuasion is to create enough doubt about the NGO's credibility that the PM eschews action under a wider range of circumstances than would have happened otherwise.

4.3 The Firm's Optimal Signal

Having established the conditions under which Bayesian persuasion is profitable, we turn now to deriving the firm's optimal signal. This signal will have two possible realizations: either the NGO is a "radical" or the NGO is "moderate," conditional on which the PM will update its beliefs about ζ . To solve for the optimal signal, we must identify the pair of posterior beliefs $\{\hat{\zeta}_M, \hat{\zeta}_R\}$ that induce payoff $C(E\Pi^{IP}(\zeta^0))$, then determine the Bayes-plausible distribution over these beliefs, and finally solve for the signal that generates this posterior distribution.

From the analysis leading to Proposition 4 it is clear that $\hat{\zeta}_R = \zeta^{NP}$ because the upward jump in profits occurs at ζ^{NP} and $E\Pi^{IP} = 0$ for all $\zeta \geq \zeta^{NP}$. It is also clear that there are two possibilities for $\hat{\zeta}_M$: $\hat{\zeta}_M = 0$ in case (a) when $\theta_H < 2\theta_L$ and case (b) when $\theta_H > 2\theta_L$ and $F > F_H/4$ and $\hat{\zeta}_M = \zeta^T$ in case (c) when $\theta_H > 2\theta_L$ and $F < F_H/4$. It is then straightforward to compute the distributions on these posteriors that are consistent with Bayes plausibility. In cases (a) and (b), the optimal signal must generate posterior belief ζ^{NP} with probability ζ^0/ζ^{NP} and posterior belief 0 with probability $1 - \zeta^0/\zeta^{NP}$. In case (c), the optimal signal must generate posterior belief ζ^{NP} with probability $(\zeta^0 - \zeta^T)/(\zeta^{NP} - \zeta^T)$ and posterior belief ζ^T with probability $(\zeta^{NP} - \zeta^0)/(\zeta^{NP} - \zeta^T)$.

We are now in a position to solve for the firm's optimal signal by making use of equations (9) and (10).

In cases (a) and (b), we set $\widehat{\zeta}_M = 0$ and $\widehat{\zeta}_R = \zeta^{NP}$, and in case (c) we set $\widehat{\zeta}_M = \zeta^T$ and $\widehat{\zeta}_R = \zeta^{NP}$. It is straightforward to show the following result.

Proposition 5 *Assuming $\zeta^{NP} < 1$, then if $\zeta^0 < \zeta^{NP}$ and (a) $\theta_H < 2\theta_L$ or (b) $\theta_H > 2\theta_L$ and $F > F_H/4$, the firm's optimal signal is*

$$(\pi_M, \pi_R) = \left(\frac{\zeta^0(\sqrt{F} - \bar{\theta})}{(1 - \zeta^0)q(\theta_H - \sqrt{F})}, 1 \right) \quad (14)$$

and if (c) $\zeta^0 \in (\zeta^T, \zeta^{NP})$, $\theta_H > 2\theta_L$ and $F < F_H/4$, the firm's optimal signal is

$$(\pi_M, \pi_R) = \left(\frac{(\zeta^0 - \zeta^T)(1 - \zeta^{NP})}{(\zeta^{NP} - \zeta^T)(1 - \zeta^0)}, \frac{(\zeta^0 - \zeta^T) \zeta^{NP}}{(\zeta^{NP} - \zeta^T) \zeta^0} \right)$$

or

$$\begin{aligned} \pi_M &= \frac{(\sqrt{F} - \bar{\theta})}{q(1 - q)(1 - \zeta^0)} \frac{(\zeta^0(1 - q)(2\sqrt{F} - \theta_L) - q(\theta_H - 2\sqrt{F}))}{\sqrt{F}(\theta_H - \theta_L)} \\ \pi_R &= \frac{(\theta_H - \sqrt{F})}{(1 - q)\zeta^0} \frac{(\zeta^0(1 - q)(2\sqrt{F} - \theta_L) - q(\theta_H - 2\sqrt{F}))}{\sqrt{F}(\theta_H - \theta_L)}. \end{aligned} \quad (15)$$

The proposition shows that the firm's optimal signal differs significantly from cases (a) and (b) to case (c). In cases (a) and (b), if the NGO really is a radical then the firm's signal is designed so that its realization always indicates that the NGO is a radical. This is intuitively appealing: the firm chooses a signal that does not distort the NGO's type when it is truly a radical, but sends a noisy signal when the NGO is truly a moderate, thereby increasing the PM's doubt about the credibility of the NGO. The firm's signal generates a posterior that mixes between $E\Pi^{IP}(0) = -cq(p_H)^2$ and $E\Pi^{IP}(\zeta^{NP}) = 0$. With Bayesian persuasion, the PM has enough doubt about the NGO that if the firm's "investigation" says the NGO is a radical and the NGO subsequently lobbies, then the PM does not bother setting a policy because its posterior belief that the NGO is a radical is ζ^{NP} . If the NGO fails to lobby, of course, the PM takes no policy action regardless of the realization of the firm's signal. Only if the "investigation" says the NGO is a moderate and the NGO subsequently lobbies does the PM set a policy, and in this case it sets the toughest possible policy, p_H . Thus, the firm is consciously taking a risk of receiving the worst possible payoff, which occurs when policy is set

at p_H . However, the risk of obtaining this policy is more than offset by the increased likelihood of receiving no policy at all.¹⁶

In case (c) the firm takes a more cautious approach, and even if the NGO really is a radical the firm's signal induces a posterior that places strictly positive probability on the realization indicating the NGO is a moderate. The reason is that the firm's expected profit function is not strictly decreasing in ζ , and $E\Pi^{IP}(\zeta^T) > E\Pi^{IP}(0)$, so the firm's payoff is higher if the posterior it induces mixes between ζ^T and ζ^{NP} instead of between 0 and ζ^{NP} .

It is worth pointing out that doubt creation also benefits the PM in our model. This does not normally happen in models of Bayesian persuasion because the receiver of information does not incur a cost of taking action.

Proposition 6 *If the firm engages in Bayesian persuasion, then the policy maker's expected utility is increased.*

Intuitively, the PM's fixed cost of taking action imposes a utility loss, and Bayesian persuasion reduces the probability that the PM will need to incur this expense. Without the firm's signal, the PM incurs the cost F to take action with probability $q + (1 - q)\zeta$. In cases (a) and (b), when the firm sends its optimal signal the PM only sets a policy when it is confident the state is high, which occurs with probability $q(1 - \zeta)(1 - \pi_M) < q + (1 - q)\zeta$. The reduction in expected costs of taking action thus creates benefits for the PM. In case (c), the PM is never certain the state is high because $\pi_R < 1$, so even if the PM observes the signal realization indicating the NGO is a "moderate," the PM cannot infer with certainty that this is true. Nevertheless, the PM is less likely to act than in the case without Bayesian persuasion, and thus

¹⁶This result differs sharply from Kamenica and Gentzkow's (2011) lobbying example, in which the lobbyist never benefits from the use of Bayesian persuasion. Of course, their lobbyist is trying to directly influence the PM, while our firm is trying to indirectly influence the PM by creating doubt about the expert's signal, but that is not the fundamental reason for the difference in results, as Proposition 4 above shows. Rather, it is the presence of fixed costs of implementing policy that creates the difference. These costs enable Bayesian persuasion to dissuade the PM from taking action unless it is certain that the state is high, and this is profitable for the firm.

benefits from a reduction in the expected costs of taking action. Of course, this does not imply that social welfare more broadly is increased by doubt creation strategies, but it does show that doubt creation enables a savings in policy costs that benefits the PM.

5 Implementing Doubt Creation via a Think Tank

The commitment assumption embodied in the use of Bayesian persuasion may seem a strong one, since the firm has incentives *ex post* to suppress signals that indicate the NGO is a moderate. Thus, in this section we consider an alternative method for disseminating doubt that embodies a greater amount of investment and hence may offer a more plausible means of generating commitment. Specifically, we consider the creation of a “think tank” (TT) that openly professes to be “pro-business” or “brown,” but whose bias is not known with certainty. As Rich (2005, p. 20) points out, “The number of new think tanks more than quadrupled between 1970 and 2000...More than half of the new think tanks that formed in this period were identifiably ideological. Two-thirds of these were identifiably conservative.” We show in this section that in the context of our model, the firm does not want to create a TT that is a fully credible moderate; instead, it prefers to create a “brown” TT whose bias is not known with certainty to the PM.

The desired level of uncertainty can be created by employing a mix of people whose apparent biases vary. For example, the American Enterprise Institute (AEI), one of the best-known “conservative” think tanks in the U.S., has a “Scholars” list that includes some like Kevin Hassett and Stephen Oliner, who are known for their many publications in refereed academic journals, and others like John Bolton and Paul Wolfowitz, who are known for their political experience and opinions.¹⁷ We consider explicitly the optimal design of such a think tank from the firm’s perspective, i.e. the distribution of biases that maximize profits for the firm, and whether it can implement the benchmark from our Bayesian persuasion analysis. Of course, personnel decisions are not permanent and irreversible, but the roster of fellows at any given think tank tends to change relatively slowly, thus engendering a degree of commitment. Moreover, few think tanks speak with

¹⁷See <http://www.aei.org/scholars/#ae>, accessed on June 28, 2017.

one entirely coherent voice, especially those that seek to attract fellows with some degree of intellectual credibility. Thus, funding a think tank provides a potential commitment mechanism for firms seeking to influence the public policy process.

Consider the issue of climate change. AEI’s Kevin Hassett (on leave from 2017-2019 to serve as chairman of the Council of Economic Advisors) has written in favor of carbon taxes,¹⁸ a mainstream position among environmental economists. However, Peter Wallison, who received a law degree in 1966 and worked on financial deregulation in the Reagan Administration, is also a fellow at AEI and takes a more extreme view, arguing that climate activists are in the thrall of a “secular religion and are not really concerned with human welfare.”¹⁹ Benjamin Zycher, who holds a PhD in economics but has never published a peer-reviewed paper on environmental economics, is the John Searle Chair and Resident Scholar at AEI, with stated expertise in energy and environmental economics.²⁰ On climate change, he writes “By failing to withdraw from the [Paris] agreement, the Trump administration is shooting itself, and all Americans, in the foot.”²¹ It is hard to tell from this mix of positions where AEI stands on the issue of climate change or appropriate policy responses to it. Thus, by choosing a mix of fellows with varying views on an issue, a think tank can associate itself with a randomized set of positions, and thereby establish a probability of being a radical or a moderate.²²

5.1 Basic Model Structure

Following our earlier treatment of the green NGO, we assume the TT can have two possible degrees of bias.

It may be a radical “brown” with probability γ , or a moderate brown with probability $1 - \gamma$. As with

¹⁸<https://www.aei.org/publication/climate-change-caps-vs-taxes/>

¹⁹For his comments, see <http://www.aei.org/publication/the-anti-climate-change-movement-is-not-really-about-climate-change/> For his background, see <http://www.aei.org/wp-content/uploads/2009/01/Wallison-Peter-Biography-1.pdf> Both were accessed November 29, 2017.

²⁰<http://www.aei.org/wp-content/uploads/2014/04/Zycher-cv-March-2017.pdf> Accessed November 29, 2017.

²¹<http://www.aei.org/tag/climate-change/> Accessed November 29, 2017.

²²We are assuming that reputation rests with the think tank itself, rather than with the individual scholars. In future work it may be of interest to delve further into this assumption, but it does seem consistent with the media’s common habit of reporting on “a Harvard study” without mentioning the authors’ names.

our analysis of the NGO, a radical TT cannot credibly convey its information to the decisionmaker, but a moderate TT can do so.²³

If the TT is a radical brown, then it will always lobby and claim that the state is θ_L , regardless of the value of θ it observes. However, from the PM's perspective, when the TT lobbies, the TT might be a radical brown or it might be a moderate brown that in fact observed θ_L . In contrast, if the TT observes the state is $\theta = \theta_H$, and thus does not lobby, the PM can be sure the state is θ_H .

The PM now receives two signals about the state of the world. Each lobbyist observes the true state perfectly, but the PM is unsure of their biases so is uncertain what to infer when lobbying occurs. Note that there is never a state of the world in which neither party lobbies, so there are three possible things the PM may observe:

1. TT lobbies and NGO lobbies. Conditional on the state being low, this happens with probability ζ . Conditional on the state being high, this occurs with probability γ . Thus, the total probability that both players lobby is $(1 - q)\zeta + q\gamma$. When both players lobby the PM infers

$$\omega = \Pr(\theta = \theta_H | \text{Both players appear}) = \frac{\gamma q}{(1 - q)\zeta + \gamma q}.$$

Given this updated belief, the PM's optimal policy is to set $p_\omega = \omega\theta_H + (1 - \omega)\theta_L$. It is straightforward to show that if $0 < \gamma < 1$ and $0 < \zeta < 1$, then $\omega < \rho$ and hence $p_\omega < p_\rho$.

2. TT lobbies and NGO does not lobby. This can only happen if the state is low. Conditional on the state being low, the NGO does not lobby with probability $1 - \zeta$. The total probability of this case is $(1 - q)(1 - \zeta)$. If the PM observes this case, it can infer with certainty that the state is low.
3. TT does not lobby and NGO does lobby. This can only happen if the state is high. Conditional on the state being high, the TT does not lobby with probability $1 - \gamma$. The total probability of this case is $q(1 - \gamma)$. If the PM observes this case, it can infer with certainty that the state is high.

²³The analysis of the range of lobbying costs for which the moderate TT can credibly lobby closely parallels that for the NGO, so we do not present it here.

The existence of the TT has two effects. First, when the PM is uncertain of the state (because both parties lobbied), it sets a weaker policy than it would have otherwise (We have shown that $p_\omega < p_\rho$). Second, the PM now has a chance of being certain that the state is high, in which case it sets a policy of p_H . (Note that the probability the PM is able to set the correct policy when the state is low remains unchanged.)

As before, we assume $F \in (F_L, F_\rho)$. With a single noisy NGO, the PM will take action when the NGO lobbies, and expected profits are

$$E\Pi^{IP} = -c[q + (1 - q)\zeta]p_\rho^2.$$

By comparison, expected profits with lobbying by the noisy NGO and the noisy TT are

$$E\Pi^{TT} = -cq(1 - \gamma)p_H^2 - c[q\gamma + (1 - q)\zeta]p_\omega^2. \quad (16)$$

Note that the TT is strictly worse for the firm than just having the NGO—*unless* the TT can create enough doubt that the PM will not act when both parties lobby.²⁴ The PM's net benefit if it takes action in this situation is just as in the noisy NGO case except that the relevant belief is ω instead of ρ :

$$EU_M^{TT} - EU_M^{NP} = (\omega p_H + (1 - \omega)p_L)^2 - F > 0.$$

The PM will act if and only if

$$\omega = \frac{\gamma q}{(1 - q)\zeta + \gamma q} > \rho^{NP} = \frac{\sqrt{F} - \theta_L}{(\theta_H - \theta_L)}. \quad (17)$$

It is now possible to characterize the optimal design of the think tank: clearly the firm would like to keep γ as large as possible, subject to ensuring that $\omega \leq \rho^{NP}$.

Proposition 7 *For a firm supporting a think tank, the optimal probability that the think tank is a radical is*

$$\gamma^* = \frac{(1 - q)\zeta(\sqrt{F} - \theta_L)}{q(\theta_H - \sqrt{F})} = \frac{\zeta}{\zeta^{NP}}. \quad (18)$$

The firm's expected profits with the optimal TT strategy are

$$E\Pi^{TT} = -cq(1 - \gamma^*)p_H^2. \quad (19)$$

²⁴It is straightforward to show that expected profits with the TT are strictly concave, and that expected profits are increasing in γ at $\gamma = 1$. This parallels our earlier result that Bayesian persuasion is not profitable when $F < F_L$.

It is straightforward to derive the comparative statics of the optimal TT bias, as summarized in the following proposition.

Proposition 8 *The firm's expected profits are maximized when the think tank has a probability $\gamma^* = \zeta/\zeta^{NP}$ of being a radical. This probability is strictly increasing in ζ and F , and strictly decreasing in q . In addition, the rate at which γ^* increases with ζ is decreasing in q and increasing in F . Assuming Bayesian persuasion is profitable, the firm prefers for the TT to be more radical than the NGO.*


The results of the proposition are shown visually in Figure 4. The bold solid lines show the positive relationship between γ^* and ζ for different values of q and F , indicating that the TT becomes more likely to be an extremist as the NGO becomes more likely to be an extremist. In addition, the line tracing out γ^* as a function of ζ rotates clockwise with increases in q and counterclockwise with increases in F , as shown by the heavy dashed curve. Thus, when a severe state of the world becomes more likely, the TT adopts a more moderate posture in order to have a greater chance of influencing policy. On the other hand, when the PM's cost of action increases, the TT increases its radical tendencies to increase the chance that the PM decides not to take any action. In addition, we see that when Bayesian persuasion is profitable (or equivalently, $F > \bar{\theta}^2$, so the PM would not take action when it is uninformed) the TT prefers to be more radical than the NGO, as in the shaded upper region of Figure 4.

5.2 Implementing Bayesian Persuasion via a Think Tank

In this section we address the conditions under which the think tank can implement the BP benchmark. The next proposition examines whether the BP benchmark can be implemented by means of a think tank.

Proposition 9 *Assuming $\zeta^0 < \zeta^{NP} < 1$, then if (a) $\theta_H < 2\theta_L$ or (b) $\theta_H > 2\theta_L$ and $F > F_H/4$, a think tank can induce the firm's optimal distribution over policy outcomes, and hence can be used to implement the Bayesian persuasion benchmark of doubt creation. However, if (c) $\zeta^T < \zeta^0 < \zeta^{NP} < 1$, $\theta_H > 2\theta_L$*

Title



4.png

Figure 4: Optimal Probability of Extremism for the Think Tank

and $F > F_H/4$, then a think tank is profitable for a smaller range of ζ^0 than is Bayesian persuasion, and it cannot implement the Bayesian persuasion benchmark.

Thus, the creation of a think tank can be an effective tool for implementing a policy of doubt creation under certain circumstances. In some cases, a think tank can implement the Bayesian persuasion benchmark. In other cases, however, even though the think tank is profitable it cannot achieve the Bayesian persuasion benchmark; it may not even be profitable at all, even though the benchmark would be. These possibilities are illustrated in Figure 5, which shows case (c) with $\zeta^T < \zeta^0 < \zeta^{NP} < 1$, $\theta_H > 2\theta_L$ and $F > F_H/4$. The concave closure created by Bayesian persuasion is shown with bold dashes, and the expected profit curve for the Think Tank is shown with bold dots. It is clear from the construction of the two curves that the expected profits from the Think Tank always fall below the Bayesian persuasion benchmark. Moreover, it is clear that the range of values for which the Think Tank is profitable, $(\hat{\zeta}, \zeta^{NP})$ is smaller than the range (ζ^T, ζ^{NP}) for which Bayesian persuasion is profitable. Because employment decisions may create a greater degree of commitment than simply promising to reveal the results of a persuasive “experiment,” the think tank may be an attractive tool for a firm seeking to delay or block regulation. Note that under the conditions

5 TT vs BP




Figure 5: Expected Profits from Bayesian Persuasion vs. a Think Tank in Case (c)

of the proposition, the TT is operating in the upper shaded region of Figure 4, and hence prefers to be more radical than the NGO.

It is straightforward to show that doubt creation is more profitable when the PM has higher costs of taking action, and it is intuitive since this makes it easier to achieve a level of doubt sufficient to block action. It is also easy to show that doubt creation is more beneficial when the NGO is widely believed to be trustworthy, as this is precisely the case when the NGO may convince the PM to implement a stringent policy.

6 Conclusions

In this article, we have extended a workhorse model of political decision-making to incorporate the policy-maker’s fixed costs of taking action, and shown how the creation of doubt can be profitable for a firm under these circumstances. In this model, the policy maker is uncertain of the state of the world, and relies on information from an NGO to shape its decisions. We considered two strategies the firm can use. First, the firm can create doubt about the credibility of the NGO. Second, the firm can create a “brown” think

tank that serves as a lobbyist, but whose bias is not known with certainty by the PM. We showed that both strategies can implement the same optimal level of doubt, and we discussed how they can be implemented. By construction these strategies are profitable for the firm. More surprisingly, they are also beneficial for the policy maker because they reduce its expected cost of taking action.

There are many related ideas that are worthy of further research. It would be worthwhile to study further the possibility that policy is not promulgated, and that multiple NGOs then have the opportunity to engage in private politics along the lines of Baron (2012). This would allow for analysis of the conditions under which the NGO prefers public politics to private politics, and vice versa. There is also a need to embed NGO strategic choices about lobbying within a model of competition between groups for membership and funding, such as that of Heyes and Martin (2017).

The parallels between the strategies identified here and the use of “astroturf” lobbying groups are striking and worthy of further study. Walker (2014) provides a richly fleshed out account of exactly how these groups achieve their objectives. Lyon and Maxwell (2004) present a formal theory of astroturf groups in which firms covertly fund grassroots lobbying groups in order to influence legislative decisions in their favor. This strategy implicitly creates doubt about whether grassroots groups are legitimate or simply fronts for corporate interests. The strategy examined in the present paper involves doubt about a lobbyist’s objectives rather than its funding sources, as emphasized by Leitzinger and Terlaak (2013). Interestingly, Leitzinger (2014) uses uncertainty on the two dimensions of funding and objectives to classify a wide range of corporate non-market strategies ranging from greenwash to astroturf to corporate social responsibility to think tanks. A formal model integrating both types of uncertainty might be a valuable contribution to extending non-market strategy research.

Think tanks are in particular need of more research. Their numbers have increased rapidly in the last 20 years (Djelic, 2014), but their impacts on policy and their use as part of corporate non-market strategy are only beginning to be studied. Barley (2010) notes that during the 1970s and 1980s the number of new conservative think tanks created was more than twice the number of new liberal think tanks. As explained by Jones (2014), this growth in conservative think tanks was part of a larger ideological movement

commonly referred to as neoliberalism, which placed great importance on the value of unregulated free markets. Consistent with our analysis, Barley (2010, p. 792) points out that “the newcomers were less likely than their predecessors to seek reputations for unbiased objective analysis.” Brulle (2014) identifies 91 organizations that lobby against climate action and together have an annual income of just over \$900 million; the largest source of funds was a group called Donors Trust whose contributors are kept anonymous, and the recipient of the most funding was the American Enterprise Institute. If it is valuable for such think tanks to be publicly unclear about their political bias, as we find, how exactly should they go about creating such doubt? One natural approach is to employ a mix of “fellows,” some of whom have solid academic credentials and others of whom are simply advocates for certain political positions. This is exactly what Barley (2010, p. 792) observes: “Like the older think tanks, the newcomers employed researchers to conduct studies and write articles and monographs, but they were less likely to write books or publish in established scholarly journals.” Given the growing importance of think tanks, more theoretical and empirical research into their role in society seems highly worthwhile.

7 Appendix

Proof of Proposition 1: Consider first the case when $F < F_L$. In an informative equilibrium, the PM will set $p_L = \theta_L$ when the NGO does not lobby and $p_H = \theta_H$ when the NGO does lobby. We must show that it is indeed optimal for the NGO to lobby if and only if the state is high. When the state is θ_L , if the NGO does not lobby it will receive payoff $-\delta^2$ and if it does lobby it will receive payoff $-(\theta_H - \theta_L - \delta)^2 - L$. Some algebra shows that not lobbying is preferred if

$$\delta < \bar{\delta} \equiv \frac{L}{2(\theta_H - \theta_L)} + \frac{(\theta_H - \theta_L)}{2}.$$

When the state is θ_H , if the NGO does not lobby it will receive payoff $-(\theta_L - \theta_H - \delta)^2$ and if it does lobby it will receive payoff $-\delta^2 - L$. Some algebra shows that lobbying is preferred if

$$\delta > \underline{\delta} \equiv \frac{L}{2(\theta_H - \theta_L)} - \frac{(\theta_H - \theta_L)}{2}.$$

Thus, when $F < F_L$ an informative equilibrium exists for $\delta \in (\underline{\delta}, \bar{\delta})$, with $\bar{\delta} - \underline{\delta} = \theta_H - \theta_L$.

Consider now the case when $F \in (\theta_L, \theta_H)$. Since $F > F_L$, the PM will not act if the state is θ_L so there is no reason for the NGO to lobby in this state. If the state is θ_H but the NGO does not lobby, the PM will infer that the state is θ_L and choose to take no action, giving the NGO utility of $-(\theta_H + \delta)^2$. If the state is θ_H and the NGO does lobby, the PM sets $p_H = \theta_H$ and the NGO obtains utility of $-\delta^2 - L$. The NGO prefers to lobby in state H if $-\delta^2 - L > -(\theta_H + \delta)^2$, or

$$\delta > \underline{\delta} \equiv \frac{L}{2\theta_H} + \frac{\theta_H}{2} - \theta_H.$$

If the state is θ_L and the NGO does not lobby, the PM will infer that the state is θ_L and choose to take no action, giving the NGO utility of $-(\theta_L + \delta)^2$. If the state is θ_L and the NGO does lobby, the PM sets $p_H = \theta_H$ and the NGO obtains utility of $-(\theta_H - \theta_L - \delta)^2 - L$. The NGO prefers not to lobby in state θ_L if $-(\theta_L + \delta)^2 > -(\theta_H - \theta_L - \delta)^2 - L$ or

$$\delta < \bar{\delta} \equiv \frac{L}{2\theta_H} + \frac{\theta_H}{2} - \theta_L.$$

Thus, when $F \in [F_L, F_H)$, an informative equilibrium exists for $\delta \in (\underline{\delta}, \bar{\delta})$, with $\bar{\delta} - \underline{\delta} = \theta_H - \theta_L$. Thus the range of biases that support an informative equilibrium is of exactly the same magnitude in the two cases. Some algebra shows that $\underline{\delta} < \underline{\delta} < \bar{\delta} < \bar{\delta}$. ■

Proof of Lemma 3: When the NGO lobbies, the PM updates its belief to $\rho = q/(q + (1 - q)\zeta)$ and the PM's expected utility of taking an imperfectly informed policy, conditional on NGO lobbying, is

$$EU_M^{IP} \Big|_{NGO \text{ lobbies}} = -(1 - \rho)\rho(p_H - p_L)^2 - F.$$

Alternatively, the PM could opt to take no action and obtain

$$EU_M^{NP} \Big|_{NGO \text{ lobbies}} = -(1 - \rho)(p_L)^2 - \rho(p_H)^2.$$

Thus, when the NGO lobbies, the PM acts if

$$EU_M^{IP}|_{NGO \text{ lobbies}} - EU_M^{NP}|_{NGO \text{ lobbies}} = (\rho p_H + (1 - \rho)p_L)^2 - F > 0. \quad (20)$$

The expression $EU_M^{IP} - EU_M^{NP} = (\rho p_H + (1 - \rho)p_L)^2 - F > 0$ can be solved using the quadratic formula to obtain an expression for ρ^{NP} , which is easily converted to an expression in terms of ζ . From this perspective, the PM takes action only if

$$\zeta < \zeta^{NP} \equiv \frac{q}{1 - q} \frac{(p_H - \sqrt{F})}{(\sqrt{F} - \theta_L)}.$$

■

Proof of Proposition 6: Begin with the reference case without Bayesian persuasion. In this case, the NGO lobbies with probability $q + (1 - q)\zeta$, in which case the PM gets $-(1 - \rho)\rho(\theta_H - \theta_L)^2 - F$. The NGO does not lobby with probability $(1 - q)(1 - \zeta)$ in which case the PM gets $-\theta_L^2$. The PM's expected utility is thus

$$EU_M^{IP} = -q(1 - \rho)(\theta_H - \theta_L)^2 - (1 - q)(1 - \zeta^0)\theta_L^2 - \frac{q}{\rho}F. \quad (21)$$

With Bayesian persuasion, we must consider both cases from Propositions 4 and 5. Consider first case (a), with $\theta_H < 2\theta_L$ or case (b), with $\theta_H > 2\theta_L$ and $F > F_H/4$, and $\zeta^0 < \zeta^{NP} < 1$. The PM takes no action if its posterior is ζ^{NP} . The only time the PM takes action is when it is certain the state is high, which occurs with probability $q(1 - \lambda) = q(1 - \pi_M)(1 - \zeta^0)$. Expected welfare is thus

$$EU_M^{BP} = -q(1 - \lambda)F - q\lambda\theta_H^2 - (1 - q)\theta_L^2.$$

Some algebra shows that Bayesian persuasion increases the PM's utility because

$$\Delta EU = \frac{q(1 - q)\zeta^0}{q + (1 - q)\zeta^0}(\theta_H - \theta_L)^2 - (1 - q)\zeta^0\theta_L^2 - q\lambda\theta_H^2 + Fq\left(\frac{1}{\rho} - (1 - \lambda)\right) > 0.$$

Consider now case (c), with $\zeta^{NP} < 1$, $\theta_H > 2\theta_L$ and $F < F_H/4$ and $\zeta^0 \in (\zeta^T, \zeta^{NP})$. Without Bayesian persuasion, equation (21) still expresses the PM's expected utility. With Bayesian persuasion, the PM again takes no action if its posterior is ζ^{NP} . However, in case (c) the PM is never certain the state is high because $\pi_R < 1$, so even if the PM observes the signal realization indicating the NGO is a "moderate," the PM

cannot infer with certainty that this is true. Thus the PM's expected utility with Bayesian persuasion is of the same general form as (21) but using ζ^T instead of ζ^0 , and multiplied by $1 - \lambda$, the probability the PM receives the signal realization "moderate". Thus we have:

$$EU_M^{BP} = (1 - \lambda) \left[-q(1 - \rho^T)(\theta_H - \theta_L)^2 - (1 - q)(1 - \zeta^T)\theta_L^2 - \frac{q}{\rho^T}F \right].$$

Note that $\zeta^T < \zeta^0$ and $\rho^T > \rho^0$. We seek to sign

$$\begin{aligned} EU_M^{BP} - EU_M^{IP} &= -q(\theta_H - \theta_L)^2[(1 - \lambda)(1 - \rho^T) - (1 - \rho^0)] - (1 - q)\theta_L^2[(1 - \lambda)(1 - \zeta^T) - (1 - \zeta^0)] \\ &\quad - qF\left[\frac{(1 - \lambda)}{\rho^T} - \frac{1}{\rho^0}\right]. \end{aligned}$$

To sign this expression, we make three observations. First, $(1 - \lambda)(1 - \rho^T) - (1 - \rho^0) < (1 - \lambda)(1 - \rho^0) - (1 - \rho^0) = -\lambda(1 - \rho^0) < 0$. Second, $(1 - \lambda)(1 - \zeta^T) - (1 - \zeta^0) < (1 - \lambda)(1 - \zeta^T) - (1 - \zeta^T) = -\lambda(1 - \rho^0) < 0$. Third, $(1 - \lambda)/\rho^T - 1/\rho^0 < (1 - \lambda)/\rho^0 - 1/\rho^0 = -\lambda/\rho^0 < 0$. Thus, $EU_M^{BP} - EU_M^{IP} > 0$, and Bayesian persuasion is beneficial for the PM also in case (c). ■

Proof of Proposition 7: For a given level of ζ , the firm would like to reduce ω just to the point of blocking action, i.e. to set $\gamma q / ((1 - q)\zeta + q\gamma) = (\sqrt{F} - \theta_L) / (\theta_H - \theta_L)$, and solving for γ^* yields the result. ■

Proof of Proposition 8: Differentiating (18), we see

$$\begin{aligned} \frac{d\gamma^*}{d\zeta} &= \frac{(1 - q)(\sqrt{F} - \theta_L)}{q(\theta_H - \sqrt{F})} > 0, \\ \frac{d\gamma^*}{dF} &= \frac{(1 - q)\zeta(\theta_H - \theta_L)}{2\sqrt{F}q(\theta_H - \sqrt{F})^2} > 0, \end{aligned}$$

and

$$\frac{d\gamma^*}{dq} = \frac{-\zeta(\sqrt{F} - \theta_L)}{q^2(\theta_H - \sqrt{F})} < 0.$$

It is easy to see that $d\gamma^*/d\zeta$ is decreasing in q and increasing in F . Finally, we see that $\gamma^* > \zeta$ when

$$\frac{(1 - q)(\sqrt{F} - \theta_L)}{q(\theta_H - \sqrt{F})} > 1,$$

or

$$\sqrt{F} > q\theta_H + (1 - q)\theta_L = \bar{\theta}.$$

■

Proof of Proposition 9: Consider first cases (a) and (b) of Proposition 4, with $\theta_H < 2\theta_L$, or with $\theta_H > 2\theta_L$ and $F > F_H/4$, and $\zeta^0 < \zeta^{NP} < 1$. We have already shown that each mechanism can be used to induce a belief on the part of the PM that is just sufficient to block action unless the PM is certain the state is high. If the TT is able to block action when both lobbyists appear, expected profits with the Think Tank are $E\Pi^{TT} = -q(1 - \gamma^*)cp_H^2$, where γ^* is defined by (18), so

$$E\Pi^{TT} = - \left[\frac{q(\theta_H - \sqrt{F}) - (1 - q)\zeta(\sqrt{F} - \theta_L)}{(\theta_H - \sqrt{F})} \right] cp_H^2.$$

With Bayesian persuasion, the firm's expected profits when the optimal signal is sent are $E\Pi^{BP} = -(1 - \pi_M)(1 - \zeta)qcp_H^2$, where π_M is defined by (14), so

$$E\Pi^{BP} = - \left(\frac{q(\theta_H - \sqrt{F}) - (1 - q)\zeta(\sqrt{F} - \theta_L)}{(\theta_H - \sqrt{F})} \right) cp_H^2$$

and $E\Pi^{BP} = E\Pi^{TT}$. Because the two mechanisms can both be used to induce the optimal probability distribution over outcomes, they have identical impacts when used optimally.

Consider now case (c) of Proposition 4, with $\theta_H > 2\theta_L$ and $F < F_H/4$ and with $\zeta^0 \in (\zeta^T, \zeta^{NP})$ and $\zeta^{NP} < 1$. With Bayesian persuasion the PM only takes action when it observes a signal realization indicating the NGO is a moderate and the NGO lobbies, the joint probability of which is $(1 - \lambda)[q(1 - \zeta^0) + \zeta^0]$. The firm earns profits given by the point in Figure 4 on the concave closure $C(E\Pi^{IP}(\zeta^0))$. To compare this to the profitability of the TT, we express $E\Pi^{TT}$ in terms of an implied posterior distribution over ζ . The TT generates two possible posteriors: $\zeta = 0$ and $\zeta = \zeta^{NP}$, and we need the expected value of $\zeta = \zeta^0$. Let us denote by χ the probability that $\zeta = 0$ and by $(1 - \chi)$ the probability that $\zeta = \zeta^{NP}$. Then clearly we must have $\chi = (\zeta^{NP} - \zeta^0)/\zeta^{NP} = 1 - \gamma^*$. The TT is valuable for $\zeta \in (\widehat{\zeta}, \zeta^{NP})$ where $\widehat{\zeta}$ is defined by

$$-c[q + (1 - q)\widehat{\zeta}]p_\rho^2 = -q(1 - \gamma^*)cp_H^2. \quad (22)$$

Because we are in case (c), we know that

$$\left. \frac{\partial E\Pi^{TT}}{\partial \zeta} \right|_{\zeta=0} > 0,$$

which is sufficient to ensure that there exists a $\hat{\zeta} \in (0, \zeta^{NP})$ that solves equation (22). From Figure 5, it is clear that $E\Pi^{TT}(\zeta^0) < C(E\Pi^{IP}(\zeta^0))$ and that $\hat{\zeta} > \zeta^T$ ■

References

- [1] Austen-Smith, D., 1995, “Campaign Contributions and Access,” *American Political Science Review*, 89(3), 566-581.
- [2] Austen-Smith, D., 1998, “Allocating Access for Information and Contributions,” *Journal of Law, Economics, and Organization*, 14 (2), 277-303.
- [3] Barley, S. R., 2010, “Building an Institutional Field to Corral a Government: A Case to Set an Agenda for Organization Studies,” *Organization Studies*, 31, 777-805.
- [4] Baron, D. P. 2012, “The Industrial Organization of Private Politics,” *Quarterly Journal of Political Science*, 7, 135-174.
- [5] Bennesen, M. and S. E. Feldmann, 2006, “Informational Lobbying and Political Contributions,” *Journal of Public Economics*, 90, 631- 656.
- [6] Bernheim, B. D., and M. D. Whinston, 1986, “Menu auctions, resource allocation, and economic influence,” *The Quarterly Journal of Economics* 101(1), 1-31.
- [7] Bramoullé, Y. and C. Orset, 2018, “Manufacturing Doubt,” *Journal of Environmental Economics and Management*, 90, 111-133.
- [8] Brulle, R. J., 2014, “Institutionalizing Delay: Foundation Funding and the Creation of US Climate Change Counter-movement Organizations,” *Climatic Change*, 122, 681-694.
- [9] Cho, C. H., M. L. Martens, H. Kim and M. Rodrigue, 2011, “Astroturfing Global Warming: It Isn’t Always Greener on the Other Side of the Fence,” *Journal of Business Ethics*, 104, 571-587.

- [10] Cotton, C., 2009, "Should We Tax or Cap Political Contributions? A Lobbying Model with Policy Favors and Access," *Journal of Public Economics*, 93 (7-8), 831-842.
- [11] Crawford, V. P., and J. Sobel, 1982, "Strategic Information Transmission," *Econometrica*, 50, 1431-1451.
- [12] Dahm, M. and N. Porteiro, 2008, "Informational Lobbying Under the Shadow of Political Pressure," *Social Choice and Welfare*, 30, 531-559.
- [13] Daubanes, J., & Rochet, J. C, 2016, "A Theory of NGO Activism," MIT Center for Energy and Environmental Policy Research, Working Paper CEEPR WP, 10.
- [14] Djelic, M.-L., 2014, "Spreading Ideas to Change the World: Inventing and Institutionalizing the Neoliberal Think Tank," C. Garsten and A. Sörbom, eds, *Political Affair: Bridging Markets and Politics*, Edward Elgar Press.
- [15] Drutman, L., 2015, *The Business of America is Lobbying*, Oxford University Press.
- [16] Dunlap, R. E., and P. J. Jacques, 2013, "Climate Change Denial Books and Conservative Think Tanks: Exploring the Connection," *American Behavioral Scientist*, 57, 699-731.
- [17] Fleckinger, P., and M. Glachant, 2011, "Negotiating a Voluntary Agreement when Firms Self-regulate," *Journal of Environmental Economics and Management*, 62, 41-52.
- [18] Glazer, A., and H. McMillan, 1990, "Optimal Coalition Size when Making Proposals is Costly," *Social Choice and Welfare*, 7, 369-380.
- [19] Glazer, A., and H. McMillan, 1992 "Pricing by the Firm under Regulatory Threat," *The Quarterly Journal of Economics*, 107, 1089-1099.
- [20] Groseclose, T., and J. M. Snyder, 1996, "Buying supermajorities," *American Political Science Review*, 90(2), 303-315.
- [21] Grossman, G. and E. Helpman, 2001, *Special Interest Politics*, Cambridge, MA: The MIT Press.

- [22] Heyes, A., and S. Martin, 2017 “Social Labeling by Competing NGOs: A Model with Multiple Issues and Entry,” *Management Science*, 63, 1800-1813.
- [23] Hoggan, J. and R. Littlemore, 2009, *Climate Cover-up: The Crusade to Deny Global Warming*, Greystone Books Ltd.
- [24] Jacques P. J., Dunlap R. E., and Freeman M., 2008, “The Organization of Denial: Conservative Think Tanks and Environmental Scepticism.” *Environmental Politics*, 17, 349-385.
- [25] Jones, D. S., 2014, *Masters of the Universe: Hayek, Friedman, and the Birth of Neoliberal Politics*, Princeton University Press.
- [26] Kamenica, E., and M. Gentzkow, 2011, “Bayesian Persuasion,” *American Economic Review*, 101, 2590-2615.
- [27] Krishna, V., and J. Morgan, 2001, “A model of Expertise,” *The Quarterly Journal of Economics*, 116, 747-775.
- [28] Kroszner, R. S., and T. Stratmann, 1998, “Interest-group competition and the organization of congress: theory and evidence from financial services’ political action committees,” *American Economic Review*, 88, 1163-1187.
- [29] Leitzinger, J. M., 2014, *A Look Behind the Curtain: Firms’ Use of Extraintitutional Non-market Strategies*. Diss. The University of Wisconsin-Madison.
- [30] Leitzinger, J. M. and A. Terlaak, 2013, “Hidden Agendas: A Model Of Industry Engagement In Oppositional Collective Action,” In *Academy of Management Proceedings*, Vol. 2013: 17635.
- [31] Lohmann, S., 1995, “Information, Access, and Contributions: A Signalling Model of Lobbying,” *Public Choice*, 85 (3/4), 267-284.
- [32] Lyon, T. P., and J. W. Maxwell, 2004, “Astroturf: Interest Group Lobbying and Corporate Strategy,” *Journal of Economics & Management Strategy*, 13, 561-597.

- [33] Lyon, T. P., and J. W. Maxwell, 2011, “Greenwash: Corporate Environmental Disclosure under Threat of Audit,” *Journal of Economics and Management Strategy*, 20, 3-41.
- [34] Lyon, T. P., and A. W. Montgomery, 2013, “Tweetjacked: The Impact of Social Media on Corporate Greenwash,” *Journal of Business Ethics*, 118, 747-757.
- [35] Lyon, T. P., and A. W. Montgomery, 2015, “The Means and End of Greenwash,” *Organization & Environment*, 28, 223-249.
- [36] Martimort, D., and A. Semenov, 2008, “Ideological Uncertainty and Lobbying Competition,” *Journal of Public Economics*, 92, 456-481.
- [37] Maxwell, J. W., T. P. Lyon, and S. C. Hackett, 2000, “Self-Regulation and Social Welfare: The Political Economy of Corporate Environmentalism,” *The Journal of Law and Economics*, 43, 583-618.
- [38] McCubbins, M., R. G. Noll and B. R. Weingast, 1987, “Administrative Procedures as Instruments of Political Control,” *Journal of Law, Economics and Organization*, 2, 243-277.
- [39] Millner, A., & Ollivier, H., 2016, “Beliefs, Politics, and Environmental Policy,” *Review of Environmental Economics and Policy*, 10(2), 226-244.
- [40] Oreskes, N. and E. M. Conway, 2010, *Merchants of Doubt*. New York: Bloomsbury.
- [41] Potters, J., and F. Van Winden, 1992, “Lobbying and asymmetric information,” *Public Choice*, 74(3) , 269-292.
- [42] Prieur, F. and B. Zou, 2017, “On the Impact of Indirect Competition for Political Influence on Environmental Policy,” CREA Discussion Paper 17-16, University of Luxembourg.
- [43] Rich, A., 2005, “War of Ideas: Why Mainstream and Liberal Foundations and the Think Tanks they support are Losing in the War of Ideas in American Politics,” *Stanford Social Innovation Review*, 3(1): 18-25.

- [44] Roemer, T., 2015, "Why Do Congressmen Spend Only Half Their Time Serving Us?," *Newsweek*, July 29, <http://www.newsweek.com/why-do-congressmen-spend-only-half-their-time-serving-us-357995>.
- [45] Segerson, K., and T. J. Miceli, 1998, "Voluntary Environmental Agreements: Good or Bad News for Environmental Protection?" *Journal of Environmental Economics and Management*, 36, 109-130.
- [46] Shapiro, J. M., 2016, "Special Interests and the Media: Theory and an Application to Climate Change," *Journal of Public Economics*, 144, 91-108.
- [47] Sinclair-Desgagné, B., and E. Gozlan, 2003, "A Theory of Environmental Risk Disclosure," *Journal of Environmental Economics and Management*, 45, 377-393.
- [48] Snyder Jr, James M., 1990, "Campaign Contributions as Investments: The US House of Representatives, 1980-1986," *Journal of Political Economy*, 98(6), 1195-1227.
- [49] Stone, D. F., 2011, "A signal-jamming model of persuasion: interest group funded policy research," *Social Choice and Welfare*, 37, 397-424.
- [50] Viscusi, K., J. M. Vernon and J. E. Harrington, Jr., 1998, *Economics of Regulation and Antitrust*, Second Edition, The MIT Press.
- [51] Walker, E. T., 2014, *Grassroots for Hire: Public Affairs Consultants in American Democracy*, Cambridge University Press.
- [52] Yu, Z., 2005, "Environmental Protection: A Theory of Direct and Indirect Competition for Political Influence," *Review of Economic Studies*, 72, 269-286.