Communication via Third Parties*

Jacopo Bizzotto,[†] Eduardo Perez-Richet[‡] and Adrien Vigier[§]

March 2021

Abstract

A principal designs an information structure and chooses transfers to an agent that are contingent on the action of a receiver. The principal faces a trade-off between, on the one hand, designing an information structure maximizing non-monetary payoffs, and on the other hand, minimizing the information rent that must be conceded to the agent in order to implement the information structure which the principal designed. We examine how this trade-off shapes communication. Our model can be applied to study the relationship between, e.g.: political organizations and the public relations companies that campaign on their behalf, firms and the companies marketing their products, consultancies and the analysts they employ.

*We thank Vincent Anesi, Nemanja Antic, Olivier Compte, Georges Georgiadis, Gabriele Gratton, Frédéric Koessler, Vasiliki Skreta, Andy Zapechelnyuk and several seminar audiences for useful comments and feedback.

[†]Faculty of Social Sciences, OsloMet. Email: jacopo.bizzotto@oslomet.no.

[‡]Department of Economics, Sciences Po and CEPR. Email: eduardo.perez@sciencespo.fr.

[§]School of Economics, University of Nottingham. Email: adrien.vigier@nottingham.ac.uk.

1 Introduction

Communication often takes place through third parties that must be incentivized extrinsically. For example, public relations companies have become major actors on the political scene, oftentimes conducting entire campaigns on behalf of traditional parties, advocacy groups and special interest groups;[1] governments in autocratic (or illiberal) regimes design censorship and/or propaganda, but typically rely on the media to actually carry those out;[2] banks and consultancies employ analysts providing expert advice in their name;[3] large firms employ specialized companies to conduct marketing campaigns on their behalf. This reliance on third parties to inform voters and customers raises potential organizational issues, as third parties either have to be monitored, or appropriately incentivized to generate information as they are expected to. How is communication shaped by the agency problem which results from contracting with third parties?

To address this question, we develop a model of a principal choosing an information structure, τ^*, with a view to influencing a receiver's belief concerning an unknown state of the world. Said information structure is implemented by an agent, who may choose to shirk. If the agent shirks, the receiver's belief is organically affected by the surrounding environment, in keeping with a given information structure, $\breve{\tau}$, say. The focus of our paper is on the case in which effort is not contractible and where, instead, the principal is constrained to provide incentives to the agent through transfers contingent on the receiver's action.

In our setting, each information structure induces a joint distribution of the receiver's action and the state, that pins down the non-monetary payoff of the principal; the principal's agency cost of implementing τ^*, on the other hand, is determined by the *marginal* distributions of the receiver's action under, respectively, τ^* and $\breve{\tau}$. The closer said marginal distributions, the more the principal needs to pay the agent in order to induce effort. We show in Section 3 that when the agent is risk neutral, the problem of the principal takes a tractable form. The agent is rewarded *exclusively* for inducing the receiver to take a given action, $a^\dagger$, say, while τ^* is distorted (relative to the first-best information structure) so as to increase the probability of $a^\dagger$. The central trade-off of our analysis then takes the following form: increasing the probability of $a^\dagger$ reduces the non-monetary payoff of the principal,

[1] See, e.g., Sheingate (2016).
[2] We thank Gabriele Gratton for this example.
[3] We thank an anonymous referee for this example.

but enables him to lower the information rent left to the agent. The larger the agent's effort cost, the greater the pressure on the principal to distort the information structure so as to reduce the agent's information rent. In particular, in sharp contrast with traditional models of moral hazard, increasing effort cost sometimes ends up making the agent worse off.

Section 4 presents a general solution method for the principal's problem. In essence, the problem of the principal reduces to finding a binary splitting of the prior belief on two posteriors, $\hat{q}$ and $q^{\dagger}$, such that the principal obtains his concavified non-monetary payoff conditional on $\hat{q}$, but no more than his "raw" non-monetary payoff conditional on $q^{\dagger}$.[4] Section 5 fully solves two prominent examples illustrating our method and the implications of the aforementioned trade-off in the context of political campaigning.

In Section 5, the principal is a political organization (e.g., a political party, or a special interest group), the receiver a median voter, and the agent a public relations company (PRC) hired by the organization in order to sway public opinion in favor of a particular outcome. We first consider a setting with two states ($\mathcal{L}$ and $\mathcal{R}$) and two possible outcomes (L and R). This setting could capture a referendum, for instance. The preferences of the organization are partially aligned with those of the median voter, in the sense that, in each state, the organization and the voter agree on the best outcome (in Figure 1a, the solid line segments depict the organization's raw non-monetary payoffs). Hence, in the absence of an agency problem (i.e. when effort is contractible), the organization chooses a fully revealing information structure, inducing the voter to pick the socially optimal outcome in each state (in Figure 1a, the dashed line segment represents the resulting payoffs of the organization). We show that the central trade-off of our analysis then tilts the outcome of the vote towards the outcome that is ex-ante preferred by the organization (inducing payoffs represented by the dotted line segment). We thus find that, in this case, the organization chooses an information structure inducing the median voter to sometimes pick an outcome that is suboptimal ex post both for the voter and the organization.

We then examine a setting with three possible outcomes, L, M and R, which could for instance represent different candidates in an election. We assume that L is the median voter's preferred outcome in state $\mathcal{L}$, and R the median voter's preferred outcome in state $\mathcal{R}$. However, in the face of uncertainty, the voter prefers M to both extremes. By contrast,

[4] By "raw" non-monetary payoffs, we mean the principal's (expected) non-monetary payoffs expressed as a function of the receiver's posterior belief. The concavified non-monetary payoff function is the smallest concave function everywhere at least as large as the raw non-monetary payoff function.

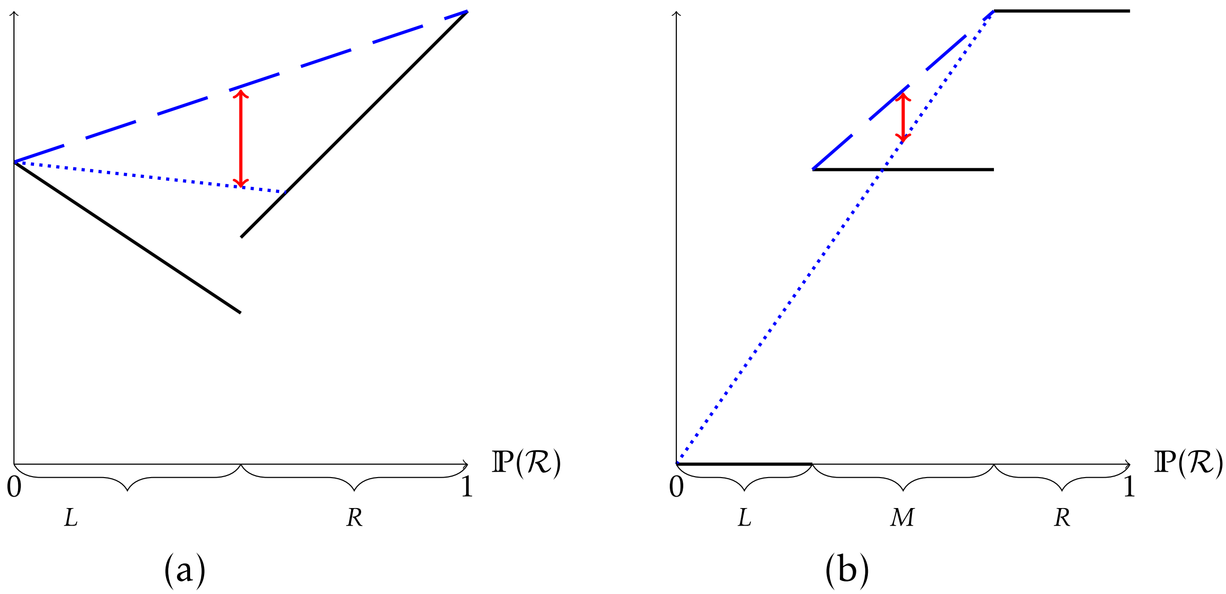

Figure 1

the organization's preferences are state-independent: L is worse than M, which is itself worse than R (in Figure 1b, the solid line segments depict the organization's raw non-monetary payoffs); but M is "closer" to R than L, meaning that in the absence of an agency problem the organization chooses an information structure inducing the voter to pick either R or M (in Figure 1b, the dashed line segment represents the resulting payoffs of the organization). We show that in this case the central trade-off of our analysis polarizes the outcome of the vote, leading the organization to choose an information structure causing the median voter to always pick one of the extremes (inducing payoffs represented by the dotted line segment). Intuitively, although M is relatively favorable to the organization, inducing it forces the organization to simultaneously lower the probability of R, which, in turn, raises the PRC's information rent.

Whether the agency problem facing the political organization helps or harms the median voter depends on the example considered: from the median voter's viewpoint, said agency problem is welfare reducing in the first example considered above, but welfare enhancing in the second example.

The special case in which the non-monetary payoffs of the principal coincide with those of the receiver is examined independently in Section 6, along with several extensions; for example, the cost of effort is allowed to depend on the choice of information structure, and the action of the receiver is allowed to have a direct impact on the welfare of the agent.

Section 7 concludes. All proofs not included in the text are in the appendix.

Related Literature. We contribute to the information design literature by exploring the implications of the observation that, in many situations, the designer might have to rely on a third party in order to implement his choice of information structure. Our paper contributes to a recent research program accounting for the fact that the design of information not only affects decisions *downstream* of information production, but also shapes incentives *upstream* of information production. In a seminal paper, Kremer, Mansour and Perry (2014) consider a designer seeking to use past consumers' experience to inform future consumers' choice. The designer would therefore like early consumers to experiment new products; however, as experimentation must be incentive compatible for the consumers, the designer is constrained in his choice of information structure. Rosar (2017) provides another important example, in which a designer can choose a test of an individual's ability, but must account for the fact that said individual could prefer not to be tested. As later studies emphasized, even when the test cannot be avoided, the kind of test chosen by the designer might influence the individual's behavior in various ways, either by affecting the individual's incentives to invest in her own ability (as in Rodina and Farragut (2017), Boleslavsky and Kim (2020) and Bizzotto and Vigier (2021)) or by affecting the individual's incentives to manipulate the test (as in Perez-Richet and Skreta (2020)), or a combination of the two (as in Kleinberg and Raghavan (2020)). Yoder (2020) is also closely connected: there, the designer contracts with a researcher that is privately informed about her cost of experimentation, and thus faces a screening problem that distorts the optimal experiment; a related problem is examined in Rappoport and Somma (2017). Finally, in Zapechelnyuk (2020), quality certification not only informs consumers but also serves to stimulate producers to supply better quality products.

Our paper is also connected to the vast literature exploring how to motivate information acquisition.[5] Most of this literature differs from our work in that the principal cannot change the technology by which agents acquire and disclose information. Two notable exceptions are Angelucci (2017), and Herresthal (2020): in the former, agents can either be required to work as a team (in order to produce a common signal), or to generate individual signals; in the latter, the agent's choice of experimentation can be made public or kept private.

[5] See, e.g., Szalay (2005), Zermeño (2011) and Carroll (2019), among many others.

Finally, our work relates to a growing literature exploring communication taking place via third parties in politics. A first strand of literature, mostly empirical, examines the changing nature of political campaigning following the advent of public relations companies (Hillygus and Shields (2009), Gibson and Rommele (2009), Strömbäck and Kiousis (2012), Tenscher, Koc-Michalska, Lilleker, Mykkänen, Walter, Findor and Jalali (2016), Grusell and Nord (2020)). A second strand of literature, theoretically studies media control in autocratic or illiberal regimes; notable recent contributions include Shadmehr and Bernhardt (2015), Boleslavsky, Shadmehr and Sonin (2020), Gratton and Lee (2020), and Kolotilin, Mylovanov and Zapechelnyuk (2021). In these papers, the focus is on censorship, i.e., situations in which autocratic regimes cancel messages ex post. Our model, on the other hand, is closer to propaganda design.

2 General Framework

The general model comprises three risk-neutral players: a principal, an agent and a receiver. The receiver chooses an action $a \in A$ affecting her payoffs and those of the principal, both of which also depend on the state of the world $\omega \in \Omega$ with probability distribution $p \in \Delta\Omega$. Let $A = \left\{a_1, \ldots, a_m\right\}$ and $\Omega = \left\{\omega_1, \ldots, \omega_n\right\}$; so m and n denote the number of possible actions and states.

Let T_x denote the set of *splittings* of the belief $x \in \Delta\Omega$, that is, the set of probability distributions on $\Delta\Omega$ averaging to x. In the spirit of Aumann, Maschler and Stearns (1995) and Kamenica and Gentzkow (2011), the principal chooses $\tau^* \in T_p$. However, the ultimate information structure depends on the agent's effort choice $e \in \{0,1\}$: either the agent exerts effort ($e = 1$), in which case the receiver's posterior belief is drawn according to τ^*, or the agent shirks ($e = 0$), in which case said posterior is drawn according to some $\check{\tau} \in T_p$ representing information generated organically by the surrounding environment. For expository simplicity, $\check{\tau}$ induces a uniformly distributed action of the receiver.[6]

Exerting effort causes disutility $c > 0$ for the agent. The central friction of the model resides in the fact that e is not contractible. The only contractible is the receiver's action a. The agent is thus incentivized through an action-contingent transfer $t^*(a)$. The players' payoffs are given by $v(a,\omega) - t^*(a)$ for the principal, $u(a,\omega)$ for the receiver, and $t^*(a) - c\mathbf{1}_{e=1}$

[6] In general, the principal may be able to exploit the non-uniform distribution induced on A when the agent shirks in order to reduce the agency cost of inducing $e = 1$. Other than this, the main insights are exactly the same as when $\check{\tau}$ induces a uniformly distributed action; see Subsection 6.1.

for the agent.[7] Throughout, we use q to denote a generic posterior belief in $\Delta\Omega$, and denote by $\overline{A} : \Delta\Omega \to \mathcal{P}(A)$ the function mapping each posterior belief to the corresponding set of receiver-optimal actions, that is,

$$\overline{A}(q) := \arg\max_{a_j \in A} \quad \sum_{i=1}^{n} q_i u(a_j, \omega_i).$$

A tuple (τ, t, α) comprising a splitting $\tau \in T_p$, a transfer scheme $t : A \to \mathbb{R}$, and an action recommendation function $a : \text{supp}\, \tau \to A$ will be referred to as a *plan*, and be represented through the shorthand notation Ξ. Three kinds of constraints will play a central role: (i) limited liability precludes negative transfers to the agent; (ii) obedience constraints ensure that the receiver chooses an optimal action regardless of the realized posterior; (iii) incentive compatibility assures that the agent prefers exerting effort over shirking. A plan Ξ satisfying (i)-(iii) will be called a *feasible* plan.

Taking a mechanism design approach, we frame the problem of the principal as that of picking one plan within the set of feasible plans, and let $\Xi^* := (\tau^*, t^*, \alpha^*)$ denote the principal's chosen plan. Hence, our focus is on the problem:

$$\max_{\Xi} \quad \sum_{q} \tau(q) \Big[\sum_{i=1}^{n} q_i v\big(\alpha(q), \omega_i\big) - t\big(\alpha(q)\big) \Big], \tag{PP$_0$}$$

$$\text{s.t.} \begin{cases} t(a_j) \geq 0, \ \forall a_j \in A, \\ \alpha(q) \in \overline{A}(q), \ \forall q \in \text{supp}\, \tau, \\ \sum_q \tau(q) t\big(\alpha(q)\big) - c \geq \frac{1}{m} \sum_j t(a_j). \end{cases}$$

Either the maximum value of (PP$_0$) is greater than the maximum expected payoff that the principal can obtain when the agent shirks, in which case the principal chooses $(\tau^*, t^*, \alpha^*) = (\tau, t, \alpha)$, where (τ, t, α) denotes a solution of (PP$_0$), or said value is less than what the principal can obtain by letting the agent shirk, and in that case the principal sets $t^*(a) = 0$ for all $a \in A$, thereby inducing the agent to shirk.

[7] In line with our lead application, we assume that the agent is unbiased and does not care about the receiver's action per se. According to Thomas B. Edsall on the New York Times blog, *The conversion of the lobbyist from a backslapper into the "man in the gray flannel suit" has already had a significant effect on American politics. The most important is that the influence industry has become unexpectedly depoliticized and disaffected.* The case in which the agent is biased introduces an additional "efficiency motive" for distorting τ^* that is independent of the model's agency problem. This case is discussed briefly in Subsection 6.3.

Notation. We record here the basic notation and terminology used throughout the analysis. Slightly abusing notation, let

$$v(q) := \max_{a_j \in \overline{A}(q)} \quad \sum_{i=1}^{n} q_i v(a_j, \omega_i), \quad \text{and} \quad \overline{\overline{A}}(q) := \arg\max_{a_j \in \overline{A}(q)} \quad \sum_{i=1}^{n} q_i v(a_j, \omega_i).$$

Thus $\overline{\overline{A}}(q)$ indicates the actions which the principal prefers among the set of receiver-optimal actions given q, and $v(q)$ the corresponding expected payoff.

We use standard notation, $\hat{v}$, to represent the concavification of the function v, and let $T_x^{\hat{v}}$ denote the set of v-concavifying splittings of x. Thus, $\tau \in T_x^{\hat{v}}$ if and only if $\tau \in T_x$ and $\sum \tau(q)v(q) = \hat{v}(x)$. The *distortion loss* $D(\Xi)$ associated with the feasible plan $\Xi = (\tau, t, \alpha)$ is defined as the difference between the first-best non-monetary payoff of the principal and the corresponding payoff given Ξ:

$$D(\Xi) := \hat{v}(p) - \sum_{q} \sum_{i=1}^{n} \tau(q) q_i v\Big(\alpha(q), \omega_i\Big).$$

3 General Analysis

We first establish the existence of a solution of the principal's problem in which the agent is rewarded exclusively for inducing the receiver to take a given action, say, $a^\dagger$. Intuitively, the agent being risk neutral, paying her only when the likelihood of $e = 1$ is maximal minimizes the principal's agency cost of inducing $e = 1$.

Lemma 1. *There exists a plan $\Xi = (\tau, t, \alpha)$ solving* (PP_0)*, and $a^\dagger \in A$, such that $t(a_j) = 0$ for all $a_j \in A \setminus \{a^\dagger\}$. Moreover, for any such Ξ and $a^\dagger$: $\mathbb{P}(a = a^\dagger \mid \Xi) \geq \mathbb{P}(a = a_j \mid \Xi)$ for all $a_j \in A$.*

Clearly, any Ξ solving (PP_0) must be such that the incentive compatibility constraint is binding; if Ξ moreover satisfies the conditions of the last lemma, then $\mathbb{P}(a = a^\dagger \mid \Xi)t(a^\dagger) - c = m^{-1}t(a^\dagger)$, giving, after simple algebra,

$$\mathbb{P}(a = a^\dagger \mid \Xi)t(a^\dagger) = c + R\Big(\mathbb{P}(a = a^\dagger \mid \Xi)\Big),$$

where

$$R(w) := \begin{cases} \frac{c}{mw-1} & \text{if } w > \frac{1}{m}, \\ +\infty & \text{if } w \leq \frac{1}{m}. \end{cases}$$

Hence, in the previous lemma, $R\Big(\mathbb{P}(a = a^\dagger \mid \Xi)\Big)$ represents the agent's information rent under the plan Ξ.

The previous remarks suggest that the principal might gain by gearing Ξ^* so as to increase the probability of a particular action in an attempt to reduce the agent's information rent. Consequently, principal-optimal plans will in general exist that comprise "suboptimal" action recommendations, in the sense of violating the condition $\alpha^*(q) \in \overline{\overline{A}}(q)$ at some posterior belief $q \in \operatorname{supp} \tau^*$. The purpose of the next lemma is to guarantee the existence of a principal-optimal plan satisfying the latter condition at all $q \in \operatorname{supp} \tau^*$.

Lemma 2. *There exists a plan (τ, t, α) solving* (PP$_0$) *such that $\alpha(q) \in \overline{\overline{A}}(q)$, for all $q \in$ supp τ.*

Henceforth, to shorten notation, let $w(\tau) := \max_q \tau(q)$. The combination of Lemmata 1 and 2 greatly simplifies the analysis, and enables us to reduce the principal's problem to

$$\max_{\tau \in T_p} \quad \sum_q \tau(q)v(q) - R\Big(w(\tau)\Big), \tag{PP$_1$}$$

as summarized in the following lemma.

Lemma 3. *Any τ solving* (PP$_1$) *also forms part of a plan Ξ solving* (PP$_0$).[8]

We can now formulate the main trade-off of our model. In what follows, let $M(w)$ denote the minimum distortion loss achievable among all feasible plans in which at least one posterior belief is generated with probability at least as large as w, that is,

$$M(w) := \min\Big\{D(\Xi) \mid \Xi = (\tau, t, \alpha) \text{ is feasible and } w(\tau) \geq w\Big\}.$$

Note that $M(w)$ is non-decreasing in w, since increasing w tightens the inequality constraint in the program defining it.

Proposition 1. *Let w solve*

$$\min_w \Big\{M(w) + R(w)\Big\}, \tag{1}$$

and consider an arbitrary splitting $\tau \in \arg\max_{\tau \in T_p} \Big\{\sum_q \tau(q)v(q) \mid w(\tau) \geq w\Big\}$. Then τ forms part of a plan Ξ solving (PP$_0$)*, and such that the agent is rewarded exclusively for inducing the receiver to take a given action $a^\dagger$ occurring with probability w.*

[8]The converse is not true. However, if τ forms part of a plan solving (PP$_0$), then $\sum_q \tau(q)v(q) = \sum_q \tilde{\tau}(q)v(q)$ for some $\tilde{\tau}$ solving (PP$_1$).

Unlike traditional models of information design, in our model, the chosen splitting τ^* affects not only the principal's payoff resulting directly from the receiver's action, but affects also the agent's information rent. The greater the weight attached by τ^* to one particular action, the smaller the rent that must be conceded to the agent in order to induce $e = 1$. The principal thus faces a trade-off, captured by (1), between minimizing the loss from distortion (relative to the first-best splitting) or the agent's information rent.

Since $R(w)$ increases with c (regardless of w), the greater c the greater the pressure on the principal to distort the chosen splitting τ^* with a view to reducing the agent's information rent. Consequently, any increase of the effort cost results in a weakly greater distortion loss. Interestingly, unlike traditional models of moral hazard, here increasing c can also *reduce* the information rent that the principal concedes to the agent. These results are summarized in the next proposition.

Proposition 2. *Raising c results in a weakly greater distortion loss, and can result in a lower information rent for the agent.*

A cautionary note is perhaps warranted. As distortions take the form of biasing τ^* in the direction of a particular action, one is led to ask whether varying effort cost induces systematic distortions of the optimal splitting. For instance, one is tempted to intuitively think that, when v is convex, increasing c should induce a mean-preserving contraction of τ^*. Roughly, the previous intuition is correct *only* as long as the payment action does not change.[9] If v is not convex, on the other hand, then increasing c can even induce a mean-preserving spread of the optimal splitting (see Subsection 5.2 for an example).

4 Solving the Principal's Problem

We now build on the analysis of the previous section to provide a general solution method for the principal's problem (PP_0). The binary splitting of the prior supported on the posterior beliefs $q^\dagger$ and $\hat{q}$ will be represented by the pair $\{q^\dagger, \hat{q}\}$. For any $\{q^\dagger, \hat{q}\}$, we then define $w(q^\dagger, \hat{q})$ by the relation

$$p = w(q^\dagger, \hat{q})q^\dagger + \left(1 - w(q^\dagger, \hat{q})\right)\hat{q}.$$

[9]For a simple example illustrating this point consider a setting with $m = n = 2$, $p = 1/4$ and $v(q) = \max\{1/8 - q, 0\}$. It is then easy to show that there exists $c_0 > 0$ such that: for $c < c_0$, τ^* splits p on $x(c)$ and 1 (where $x(c)$ increases with c), whereas for $c > c_0$, τ^* splits p on 0 and $y(c)$ (where $y(c)$ decreases with c).

Below, $\mathbf{1}_x$ denotes the splitting of $x \in \Delta\Omega$ assigning probability 1 to x; we start this section with the following preliminary result.

Lemma 4. *There exists $\overline{c}$ such that $\tau^* = \mathbf{1}_p$ for $c \geq \overline{c}$, whereas any principal-optimal splitting is informative for $c < \overline{c}$. If v is concave, then $\overline{c} = 0$.*

Henceforth, assume $c < \overline{c}$. Thus, given τ solving (PP$_1$), we can pick $q^\dagger \in \arg\max_q \tau(q)$, and let $\hat{q} := \frac{p - w(\tau)q^\dagger}{1 - w(\tau)}$. Then $\{q^\dagger, \hat{q}\}$ is a binary splitting of the prior belief, $w(\tau) = w(q^\dagger, \hat{q})$, and, by writing

$$\sum_q \tau(q)v(q) = w(\tau)v(q^\dagger) + \Big(1 - w(\tau)\Big) \sum_{q \neq q^\dagger} \frac{\tau(q)}{1 - w(\tau)} v(q),$$

we see that

$$\begin{aligned} \sum_q \tau(q)v(q) - R\Big(w(\tau)\Big) &\leq w(\tau)v(q^\dagger) + \Big(1 - w(\tau)\Big)\hat{v}(\hat{q}) - R\Big(w(\tau)\Big) \\ &= w(q^\dagger, \hat{q})v(q^\dagger) + \Big(1 - w(q^\dagger, \hat{q})\Big)\hat{v}(\hat{q}) - R\Big(w(q^\dagger, \hat{q})\Big). \end{aligned}$$

Similarly, given a binary splitting $\{q^\dagger, \hat{q}\}$ of the prior belief, let $\hat{\tau}$ denote an arbitrary splitting in $T_{\hat{q}}^{\hat{v}}$ and define $\tilde{\tau} := w(q^\dagger, \hat{q})\mathbf{1}_{q^\dagger} + \Big(1 - w(q^\dagger, \hat{q})\Big)\hat{\tau}$. Then

$$w(q^\dagger, \hat{q})v(q^\dagger) + \Big(1 - w(q^\dagger, \hat{q})\Big)\hat{v}(\hat{q}) - R\Big(w(q^\dagger, \hat{q})\Big) \leq \sum_q \tilde{\tau}(q)v(q) - R\Big(w(\tilde{\tau})\Big).$$

Combining the previous observations shows that to any τ solving (PP$_1$) is associated a binary splitting solving

$$\max_{\{q^\dagger, \hat{q}\}} \quad w(q^\dagger, \hat{q})v(q^\dagger) + \Big(1 - w(q^\dagger, \hat{q})\Big)\hat{v}(\hat{q}) - R\Big(w(q^\dagger, \hat{q})\Big). \tag{PP$_2$}$$

More importantly, we show in the next proposition that the converse is true as well: to any $\{q^\dagger, \hat{q}\}$ solving (PP$_2$) is associated at least one splitting τ solving (PP$_1$). In other words, the problem of the principal can be reduced to finding a binary splitting of the prior belief solving (PP$_2$), as summarized in the following proposition.

Proposition 3. *Assume $c < \overline{c}$. Consider $\{q^\dagger, \hat{q}\}$ solving* (PP$_2$)*. Let $\hat{\tau} \in T_{\hat{q}}^{\hat{v}}$, and $\tau = w(q^\dagger, \hat{q})\mathbf{1}_{q^\dagger} + \Big(1 - w(q^\dagger, \hat{q})\Big)\hat{\tau}$. Then τ forms part of a plan Ξ solving* (PP$_0$)*, and such that the agent is rewarded* *exclusively for inducing the receiver to take a given action $a^\dagger$ occurring with probability $w(q^\dagger, \hat{q})$.*

From Proposition 3, we immediately obtain the following procedure for solving (PP_0):

Step I: compute $\hat{v}$, the concavification of the payoff function v;

Step II: find a binary splitting $\{q^\dagger, \hat{q}\}$ of the prior belief that solves (PP_2);

Step III: find a v-concavifying splitting $\hat{\tau}$ of $\hat{q}$;

Step IV: let $\tau = w(q^\dagger, \hat{q})\mathbf{1}_{q^\dagger} + \Big(1 - w(q^\dagger, \hat{q})\Big)\hat{\tau}$, and select any action recommendation function α that satisfies $\alpha(q) \in \overline{\overline{A}}(q)$ for all $q \in \operatorname{supp} \tau$;

Step V: choose the plan $\Xi = (\tau, t, \alpha)$, where $t\Big(\alpha(q^\dagger)\Big) = mc(mw(q^\dagger, \hat{q}) - 1)^{-1}$ and $t(a_j) = 0$ for $a_j \neq \alpha(q^\dagger)$.

5 Political Campaigning

In this section, we abandon our general analysis in favor of a more narrow setting that enables us to illustrate in practice (a) the central trade-off highlighted in Section 3, (b) the procedure presented in section 4.

In this section, the role of the principal is played by some political organization (e.g., a political party, or a special interest group), and that of the receiver is played by a median voter. The agent of our model represents a public relations company (PRC), hired by the organization in order to sway public opinion in favor of one of m outcomes. An outcome could represent a particular policy, or a candidate in an election. The binary state is uniformly distributed; we let $\mathcal{L}$ and $\mathcal{R}$ stand for ω_1 and ω_2, respectively. Lastly, for expository simplicity, beliefs will here be represented by the probability attached to state $\mathcal{R}$.

5.1 The Limits of Informative Campaigning

We first consider a setting with two possible outcomes (L and R), a referendum say. The preferences of the organization are partially aligned with those of the median voter, in the sense that, in each state, the organization and the median voter agree on the best outcome. Concretely, the median voter must choose between outcomes L and R. The voter's preferred outcome is L if $q < 1/2$, and R if $q > 1/2$. Both the organization and the median voter prefer outcome L when the state is $\mathcal{L}$ and outcome R when the state is $\mathcal{R}$. However, the organization's payoffs are such that, ex ante, the organization favors outcome R. Specifically, we normalize the organization's payoffs from outcome R to be 3 in state $\mathcal{R}$ and 0 in state $\mathcal{L}$; its payoffs from outcome L are 2 in state $\mathcal{L}$ and 0 in state $\mathcal{R}$. The solid line segments in Figure 1a depict the organization's payoffs in terms of the median voter's posterior belief. The concavification of the function v is in this case given by $2+q$.

We next elicit the organization's optimal splitting by solving the reduced problem (PP_2).

Lemma 5. *If* $\{q^\dagger, \hat{q}\}$ *solves* (PP_2) *then* $\hat{q} = 0$ *and* $q^\dagger \in (1/2, 1)$.

Proof: First, note that if $\{q^\dagger, \hat{q}\}$ is such that one of the two posterior beliefs is equal to 1 and the other equal to 0, then $\{q^\dagger, \hat{q}\}$ cannot solve (PP_2). Indeed, in this case, $w(q^\dagger, \hat{q}) = 1/2$, and so $R\Big(w(q^\dagger, \hat{q})\Big) = +\infty$.

Next, we prove that any $\{q^\dagger, \hat{q}\}$ solving (PP_2) must satisfy $q^\dagger > 1/2$ and $\hat{q} < 1/2$. We will do this by showing that, for all $x \in (0, 1/2]$ and all $y \in (0, 1/2]$, where x and y are not both equal to 1/2, the splitting $\{\frac{1}{2}+x, \frac{1}{2}-y\}$ strictly dominates $\{\frac{1}{2}-x, \frac{1}{2}+y\}$ for the problem (PP_2). Note first that $w(\frac{1}{2}+x, \frac{1}{2}-y) = w(\frac{1}{2}-x, \frac{1}{2}+y) = \frac{y}{x+y} =: w$. Hence, $\{\frac{1}{2}+x, \frac{1}{2}-y\}$ strictly dominates $\{\frac{1}{2}-x, \frac{1}{2}+y\}$ for the problem (PP_2) if and only if

$$3w(\tfrac{1}{2}+x) + (1-w)(2+\tfrac{1}{2}-y) > w(2-\tfrac{1}{2}+x) + (1-w)(2+\tfrac{1}{2}+y).$$

Substituting $w = \frac{y}{x+y}$ in the previous expression yields $\frac{2xy}{x+y} > 0$.

We now show that if $\{q^\dagger, \hat{q}\}$ solves (PP_2) then either $q^\dagger = 1$ or $\hat{q} = 0$. Consider $q^\dagger \in (1/2, 1)$ and $\hat{q} \in (0, 1/2)$. Define $w := w(q^\dagger, \hat{q})$ and $y(x) := \frac{p-wx}{1-w}$. Then, by choosing ε sufficiently small, notice that $x \in (1/2, 1)$ and $y(x) \in (0, 1/2)$ for all $x \in [q^\dagger - \varepsilon, q^\dagger + \varepsilon]$. Note too that $\{x, y(x)\}$ is a binary splitting of p with $w(x, y(x)) = w$, so for all $x \in [q^\dagger - \varepsilon, q^\dagger + \varepsilon]$:

$$\begin{aligned} w(x,y(x))v(x) + \Big(1 - w(x,y(x))\Big)\hat{v}\Big(y(x)\Big) - R\Big(w(x,y(x))\Big) \\ &= 3wx + (1-w)\Big(2+y(x)\Big) - R(w) \\ &= 2wx + \frac{5}{2} - 2w - R(w). \end{aligned}$$

Hence, $\Big\{q^\dagger + \varepsilon, y(q^\dagger + \varepsilon)\Big\}$ strictly dominates $\{q^\dagger, \hat{q}\}$ for the problem (PP_2). This establishes that if $\{q^\dagger, \hat{q}\}$ solves (PP_2) then either $q^\dagger = 1$ or $\hat{q} = 0$.

Lastly, note that the alternative $q^\dagger = 1$ is ruled out, since $w(1, \hat{q}) < 1/2$ for all $\hat{q} \in (0, 1/2)$, and so $R\Big(w(1, \hat{q})\Big) = +\infty$. ■

We can now state this subsection's main result.

Proposition 4. *If $c \geq \frac{1}{2}$ then $\tau^* = \mathbf{1}_p$. Otherwise τ^* splits p on 0 and $\frac{1}{2} + x$, where $x \in (0, 1/2)$ decreases with c.*

The first-best solution is here straightforward: since in each state the organization and the median voter agree on the best outcome, the organization chooses a splitting fully revealing the state to the median voter. With agency however, the organization is unable to implement the fully revealing splitting. The reason is that, in this case, the marginal distribution of the voter's choice is the same whether or not the PRC exerts effort; so the PRC prefers shirking regardless of the transfer scheme. The organization is thus constrained to choose a biased informative structure, sometimes inducing the voter to select outcome R (which, ex ante, is preferred by the organization) in state $\mathcal{L}$ although, ex post, the organization and the median voter agree that the best outcome is then L.

In Figure 1a, the double arrow indicates the distortion loss suffered by the organization from splitting p on 0 and $1/2 + x$, where $x < 1/2$, instead of fully revealing the state to the median voter. Note that in the present setting the organization's agency problem vis-a-vis the PRC makes the median voter unambiguously *worse off*, since it induces the organization to choose a splitting generating less information than the organization would choose in the absence of an agency problem.

5.2 Partisan Politics

We now examine a setting with three possible outcomes, L, M, and R, which could for instance represent different candidates in an election. The voter's preferred outcome is L for $q < 1 - \overline{q}$, M for $q \in (1 - \overline{q}, \overline{q})$ and R for $q > \overline{q}$, where $\overline{q} \in (1/2, 1)$. However, the organization's preferences are state independent: R is the organization's most-preferred outcome, followed by M and L. We normalize the organization's payoffs to be 1 for outcome R, 0 for outcome L, and v_M for outcome M. To shorten the exposition, we assume throughout this subsection that $2(1 - \overline{q}) > v_M > (1 - \overline{q})/\overline{q}$.[10] The solid line segments of Figure 1b depict the organization's payoffs in terms of the median voter's posterior belief. The concavification of the function v is in this case given by $\frac{v_M}{1-\overline{q}} q$ if $q \leq 1 - \overline{q}$, $v_M + \left(\frac{1-v_M}{2\overline{q}-1}\right)(q - 1 + \overline{q})$ if $q \in [1 - \overline{q}, \overline{q}]$, and 1 if $q \geq \overline{q}$.

[10]The case $v_M > 2(1 - \overline{q})$ is similar, with Lemma 7 reversed, namely, in that case, for all $x \in (0, 1 - \overline{q}]$, $\{1 - \overline{q}, \overline{q} + x\}$ strictly dominates $\{\overline{q}, 1 - \overline{q} - x\}$ for the problem (PP_2). The case $V_M < (1 - \overline{q})/\overline{q}$ is less interesting, as in that case cav v is linear on $[0, \overline{q}]$.

We next elicit the organization's optimal splitting by solving the reduced problem (PP_2). We start with two lemmata, establishing that any solution $\{q^\dagger, \hat{q}\}$ of the problem (PP_2) either satisfies $q^\dagger = \overline{q}$ and $\hat{q} \in [0, 1-\overline{q}]$ (Lemma 6), or $q^\dagger > 1-\overline{q}$ and $\hat{q} \in [\overline{q}, 1]$ (Lemma 7).

Lemma 6. *If* $\{q^\dagger, \hat{q}\}$ *solves* (PP_2) *and* $q^\dagger > \hat{q}$, *then* $q^\dagger = \overline{q}$ *and* $\hat{q} \in [0, 1-\overline{q}]$.

Proof: Consider first $q^\dagger > \overline{q}$. Then, for any $\varepsilon < q^\dagger - \overline{q}$ and any $\hat{q} < 1/2$, the splitting $\{q^\dagger - \varepsilon, \hat{q}\}$ strictly dominates $\{q^\dagger, \hat{q}\}$ for the problem (PP_2), since $v(q^\dagger - \varepsilon) = v(q^\dagger) = 1 > v_M$ and $w(q^\dagger - \varepsilon, \hat{q}) > w(q^\dagger, \hat{q})$. Hence no $\{q^\dagger, \hat{q}\}$ with $q^\dagger > \overline{q}$ is a solution of (PP_2).

Next, consider $q^\dagger \in (1/2, \overline{q})$. Firstly, reasoning as in the previous paragraph shows that, for any $\hat{q} \in [0, 1-\overline{q}]$, $\{q^\dagger, \hat{q}\}$ cannot be a solution of (PP_2). Now let $\hat{q} \in (1-\overline{q}, 1/2)$, define $w := w(q^\dagger, \hat{q})$ and $y(x) := \frac{p-wx}{1-w}$. Then, by choosing ε sufficiently small, notice that $x \in (1/2, \overline{q})$ and $y(x) \in (1-\overline{q}, 1/2)$ for all $x \in [q^\dagger - \varepsilon, q^\dagger + \varepsilon]$. Note too that $\{x, y(x)\}$ is a binary splitting of p with $w(x, y(x)) = w$, so for all $x \in [q^\dagger - \varepsilon, q^\dagger + \varepsilon]$:

$$
\begin{aligned}
&w(x,y(x))v(x) + \Big(1 - w(x,y(x))\Big)\hat{v}\Big(y(x)\Big) - R\Big(w(x,y(x))\Big) \\
&\qquad = wv_M + (1-w)\Big[v_M + \Big(\frac{1-v_M}{2\overline{q}-1}\Big)(y(x) - 1 + \overline{q})\Big] - R(w) \\
&\qquad = v_M + \Big(\frac{1-v_M}{2\overline{q}-1}\Big)\Big[p - wx - (1-\overline{q})(1-w)\Big] - R(w).
\end{aligned}
$$

Hence, $\Big\{q^\dagger - \varepsilon, y(q^\dagger - \varepsilon)\Big\}$ strictly dominates $\{q^\dagger, \hat{q}\}$ for the problem (PP_2). Combining the remarks made in this paragraph shows that no $\{q^\dagger, \hat{q}\}$ with $q^\dagger \in (1/2, \overline{q})$ is a solution of (PP_2). It ensues that if $\{q^\dagger, \hat{q}\}$ solves (PP_2) and $q^\dagger > \hat{q}$, then $q^\dagger = \overline{q}$.

Finally, notice that, for any $\hat{q} \in (1-\overline{q}, 1/2)$, the splitting $\{\overline{q}, 1-\overline{q}\}$ strictly dominates $\{\overline{q}, \hat{q}\}$ for the problem (PP_2), since $w(\overline{q}, 1-\overline{q}) > w(\overline{q}, \hat{q})$ whereas

$$
w(\overline{q}, 1-\overline{q})v(\overline{q}) + \Big(1 - w(\overline{q}, 1-\overline{q})\Big)\hat{v}(1-\overline{q}) = w(\overline{q}, \hat{q})v(\overline{q}) + \Big(1 - w(\overline{q}, \hat{q})\Big)\hat{v}(\hat{q}).
$$

■

Lemma 7. *If* $\{q^\dagger, \hat{q}\}$ *solves* (PP_2), $q^\dagger < \hat{q}$ *and* $\{q^\dagger, \hat{q}\} \neq \{1-\overline{q}, \overline{q}\}$, *then* $q^\dagger > 1-\overline{q}$ *and* $\hat{q} \in [\overline{q}, 1]$.

Proof: Consider first $\hat{q} \in (1/2, \overline{q})$. Then, for any $q^\dagger \in [0, 1-\overline{q}]$, the splitting $\{q^\dagger, \hat{q}\}$ is strictly dominated by $\{1-\overline{q}, \overline{q}\}$ for the problem (PP_2), since $v_M > (1-\overline{q})/\overline{q}$ and $w(1-\overline{q}, \overline{q}) > w(q^\dagger, \hat{q})$. Moreover, the same arguments as those used in the proof of Lemma 6 establish that no splitting $\{q^\dagger, \hat{q}\}$ with $\hat{q} \in (1-\overline{q}, 1/2)$ is a solution of (PP_2). Thus no $\{q^\dagger, \hat{q}\}$ with $\hat{q} \in (1/2, \overline{q})$ is a solution of (PP_2).

Next, consider $\hat{q} \in [\overline{q}, 1]$. For any $q^\dagger < 1-\overline{q}$, the splitting $\{1-\overline{q}, \hat{q}\}$ strictly dominates $\{q^\dagger, \hat{q}\}$ for the problem (PP_2), since $v_M > (1-\overline{q})/\overline{q}$ and $w(1-\overline{q}, \hat{q}) > w(q^\dagger, \hat{q})$.

Combining the remarks in the previous paragraphs establishes that if $\{q^\dagger, \hat{q}\}$ solves (PP_2), $q^\dagger < \hat{q}$ and $\{q^\dagger, \hat{q}\} \neq \{1-\overline{q}, \overline{q}\}$, then $q^\dagger \geq 1-\overline{q}$ and $\hat{q} \in [\overline{q}, 1]$. We now show that $q^\dagger$ must be strictly greater than $1-\overline{q}$, by proving that for all $x \in (0, 1-\overline{q}]$, the splitting $\{1-\overline{q}, \overline{q}+x\}$ is strictly dominated by $\{\overline{q}, 1-\overline{q}-x\}$ for the problem (PP_2). Note first that

$$w(\overline{q}, 1-\overline{q}-x) = w(1-\overline{q}, \overline{q}+x) = \frac{2(\overline{q}+x)-1}{2(2\overline{q}-1)+2x} =: w.$$

Hence, given $x \in (0, 1-\overline{q}]$, $\{1-\overline{q}, \overline{q}+x\}$ is strictly dominated by $\{\overline{q}, 1-\overline{q}-x\}$ for the problem (PP_2) if and only if

$$w + (1-w)\left(\frac{1-\overline{q}-x}{1-\overline{q}}\right)v_M > w v_M + (1-w).$$

Substituting $w = \frac{2(\overline{q}+x)-1}{2(2\overline{q}-1)+2x}$ in the last expression yields $2(1-\overline{q}) > v_M$. ■

Combining Lemmata 6 and 7 yields this subsection's main result.

Proposition 5. *There exist* $0 < c_1 < c_2 < c_3$ *such that:*

- *for* $c < c_1$, τ^* *splits* p *on* $1-\overline{q}$ *and* $\overline{q}$;
- *for* $c_1 < c < c_2$, τ^* *splits* p *on* 0, $1-\overline{q}$ *and* $\overline{q}$;
- *for* $c_2 < c < c_3$, τ^* *splits* p *on* 0 *and* $\overline{q}$.

For $c > c_3$, *either* $\tau^* = \mathbf{1}_p$ *or some splitting* $\{x, y\}$ *where* $x > 1-\overline{q}$ *and* $y \geq \overline{q}$ *forms part of a plan solving* (PP_0).

The broad logic of the proposition is the following. As long as $c < c_1$ the organization chooses the first-best splitting $\{1-\overline{q}, \overline{q}\}$, concavifying v at $p = 1/2$. Outcomes M and R are then generated with 50 % probability each, and the PRC obtains an information rent equal to $2c$. At $c = c_1$, the organization's marginal gain from reducing said rent exactly compensates the distortion loss incurred when "resplitting" $1-\overline{q}$ on 0 and $\overline{q}$. Between $c = c_1$ and $c = c_2$, the organization thus gradually reduces the probability of outcome M in favor of the two extremes. In the range $c_2 < c < c_3$, outcome M never occurs: outcome R is generated with probability $(2\overline{q})^{-1}$, while outcome L is generated with probability $1-(2\overline{q})^{-1}$. For $c > c_3$, the pressure to contain the PRC's information rent is such that the organization all but gives up its most-preferred outcome, and induces M with probability equal or close to 1.[11]

In Figure 1b, the double arrow indicates the distortion loss suffered by the organization from splitting p on 0 and $\overline{q}$ instead of splitting p on $1-\overline{q}$ and $\overline{q}$. At $c = c_2$ this loss is exactly compensated by the resulting reduction of the information rent left to the PRC in order to incentivize effort. In the cost range $c_2 < c < c_3$, the organization's agency problem vis-a-vis the PRC thus causes complete polarization of the outcome. Note that in this range, said agency problem actually makes the median voter *better off*, by inducing the organization to choose a splitting generating more information than the organization would generate in the absence of an agency problem.

6 Extensions

6.1 Asymmetric Action Distribution

We simplified the exposition of the baseline model by assuming that the splitting $\breve{\tau}$ (describing the distribution of the receiver's posterior belief in case $e = 0$) induces a uniformly distributed action of the receiver. We briefly discuss here how to generalize the analysis.

The incentive compatibility constraint now takes the form $\sum_q \tau(q)t\Big(\alpha(q)\Big)-c \geq \sum_q \breve{\tau}(q)t\Big(\alpha(q)\Big)$. Since the agent is risk neutral, the principal can always minimize the cost of inducing $e = 1$ by rewarding the agent exclusively for inducing the receiver to take a given action, $a^\dagger$, say. However, with asymmetric action distribution under $e = 0$, the action maximizing the likelihood that $e = 1$ need no longer be the most probable action under $e = 1$: the relevant variable is now $\mathbb{P}(a = a^\dagger \mid \Xi)/\mathbb{P}(a = a^\dagger \mid \breve{\tau})$. Hence, to solve the principal's problem we must now define one information rent function per possible action $a_j \in A$:

$$R_{a_j}(w) := \begin{cases} \frac{c\mathbb{P}(a=a^\dagger|\breve{\tau})}{w-\mathbb{P}(a=a^\dagger|\breve{\tau})} & \text{if } w > \mathbb{P}(a = a^\dagger \mid \breve{\tau}), \\ +\infty & \text{if } w \leq \mathbb{P}(a = a^\dagger \mid \breve{\tau}). \end{cases}$$

The reduced problem of the principal then becomes

$$\max_{\{q^\dagger,\hat{q}\},a^\dagger\in\overline{A}(q^\dagger)} \quad w(q^\dagger,\hat{q})v(q^\dagger) + \Big(1-w(q^\dagger,\hat{q})\Big)\hat{v}(\hat{q}) - R_{a^\dagger}\Big(w(q^\dagger,\hat{q})\Big).$$

[11] Depending on the parameters, $c_3 = \overline{c}$ or $c_3 < \overline{c}$.

If $a^\dagger$ and $\{q^\dagger, \hat{q}\}$ comprise a solution of the previous problem then, for any $\hat{\tau} \in T_{\hat{q}}^{\hat{v}}$, the splitting $\tau = w(q^\dagger, \hat{q})\mathbf{1}_{q^\dagger} + \left(1 - w(q^\dagger, \hat{q})\right)\hat{\tau}$ forms part of a plan Ξ solving the principal's initial problem. Note that if there exists $a^\dagger$ such that $\mathbb{P}(a = a^\dagger \mid \check{\tau}) = 0$, then the first-best splitting forms part of a plan Ξ solving the principal's initial problem: the agent is paid $c/\mathbb{P}(a = a^\dagger \mid \tau^*)$ for inducing the receiver to take the action $a^\dagger$, and obtains no information rent.

6.2 Full Alignment of the Principal with the Receiver

In this subsection we examine in more details the case in which $v = u$, that is, in which the non-monetary payoffs of the principal coincide with those of the receiver. For instance, the principal might be a benevolent political organization, whose interests fully align with those of the median voter. We show that in this case the problem of the principal drastically simplifies.

For each $\omega_i \in \Omega$, pick $a_{\omega_i} \in \arg\max_{a_j \in A} v(a_j, \omega_i)$. Then, as $v = u$:

$$v(q) = \max_{a_j \in A} \quad \sum_{i=1}^{n} q_i v(a_j, \omega_i). \tag{2}$$

and

$$\hat{v}(q) = \sum_{i=1}^{n} q_i v(a_{\omega_i}, \omega_i). \tag{3}$$

Plugging (2) and (3) into (PP_2) and changing the maximization variables, this reduced problem becomes

$$\max_{\substack{q^\dagger \in \Delta\Omega, \\ w \in [0,1], \\ a^\dagger \in A}} \quad \sum_{i=1}^{n} w q_i^\dagger v(a^\dagger, \omega_i) + \sum_{i=1}^{n} (p_i - w q_i^\dagger) v(a_{\omega_i}, \omega_i) - R(w),$$

and, changing the maximization variables once more,

$$\max_{\substack{x\in\mathbb{R}^n,\\ w\in[0,1],\\ a^\dagger\in A}} \quad \sum_{i=1}^n x_i v(a^\dagger,\omega_i) + \sum_{i=1}^n (p_i - x_i) v(a_{\omega_i},\omega_i) - R(w),$$

$$\text{s.t.} \quad \begin{cases} \sum_\omega x(\omega) = p, \\ 0 \le x \le \mu_0. \end{cases}$$

We thus obtain a simple linear program in x. Passing to the dual gives

$$\max_{\substack{w\in[0,1],\\ a^\dagger\in A}} \min_\lambda \quad \sum_{i=1}^n p_i \Big[\lambda - \ell(a^\dagger,\omega_i)\Big]^+ + \hat{v}(p) - w\lambda - R(w),$$

where $\ell(a^\dagger,\omega_i) := v(a_{\omega_i},\omega_i) - v(a^\dagger,\omega_i)$ represents the principal's loss from action $a^\dagger$ in state ω_i. Notice that the maximand is convex in λ and concave in w; hence, by virtue of the minimax Theorem, we can switch the order of the minimization over λ and the maximization over w. By doing this, we obtain a straightforward optimization problem in w, whose first-order condition is $\lambda + R'(w) = 0$, giving $w = m^{-1} + \sqrt{\frac{c}{m\lambda}}$. Intuitively, increasing the probability w of the action $a^\dagger$ at which the agent is rewarded lowers the latter's information rent, but implies increasing the distortion loss at a rate equal to the dual variable λ representing the shadow cost of distortion. Finally, substituting $w = m^{-1} + \sqrt{\frac{c}{m\lambda}}$ in the previous program, we obtain

$$\max_{a^\dagger\in A} \min_\lambda \quad \sum_{i=1}^n p_i \Big[\lambda - \ell(a^\dagger,\omega_i)\Big]^+ + \hat{v}(p) - \lambda m^{-1} - 2\sqrt{\frac{c\lambda}{m}}. \tag{4}$$

If $(a^\dagger,\lambda)$ solve (4), we can now retrieve a solution of (PP_2) by simply letting $w(q^\dagger,\hat{q}) = m^{-1} + \sqrt{\frac{c}{m\lambda}}$ and $q^\dagger = x/w(q^\dagger,\hat{q})$, where $0 \le x \le p$, $\sum_{i=1}^n x_i = p$, and

$$x_i = \begin{cases} 0 & \text{if } \lambda < \ell(a^\dagger,\omega_i), \\ p_i & \text{if } \lambda > \ell(a^\dagger,\omega_i). \end{cases}$$

6.3 Biased Agent

The baseline model assumed that the receiver's action affects the agent's welfare only through the payment of the principal. In this subsection, we briefly consider the case in which the action of the receiver has a direct impact on the welfare of the agent.

Let $b(a)$ represent the agent's bias towards action a, so that the agent's payoffs are now given by $t^*(a) - c\mathbf{1}_{e=1} + b(a)$. Let

$$V_b(q,w) := \max_{a_j \in \overline{A}(q)} \quad \sum_{i=1}^{n} q_i v(a_j, \omega_i) + \frac{b(a_j)}{mw-1},$$

with $V_b(q,w) = +\infty$ for $w \leq 1/m$, and

$$R_b(w) := \begin{cases} \frac{c+\frac{1}{m}\sum_{j=1}^{m} b(a_j)}{mw-1} & \text{if } w > \frac{1}{m}, \\ +\infty & \text{if } w \leq \frac{1}{m}. \end{cases}$$

The reduced problem of the principal becomes in this case

$$\max_{\{q^\dagger,\hat{q}\}} \quad w(q^\dagger,\hat{q})V_b\Big(q^\dagger, w(q^\dagger,\hat{q})\Big) + \Big(1 - w(q^\dagger,\hat{q})\Big)\hat{V}_b\Big(\hat{q}, w(q^\dagger,\hat{q})\Big) - R_b\Big(w(q^\dagger,\hat{q})\Big),$$

where $\hat{V}_b(\cdot,w)$ denotes the concavification of the function $V_b(\cdot,w)$. Intuitively, the agent's bias introduces an additional *efficiency motive* for distorting τ^* relative to the first-best splitting. However, note that, unlike the "informational rent motive" on which our paper focuses, said efficiency motive is independent of the underlying agency problem, and would still exist if effort were contractible.

6.4 Splitting-Dependent Effort Cost

In this subsection we show how the model can be adapted to incorporate a splitting-dependent effort cost.

Fix an arbitrary interior belief $\tilde{p} \in \Delta\Omega$ and let $z : \Delta\Omega \to \Delta\Omega$ map posterior beliefs of the receiver to posterior beliefs of an observer whose prior belief is $\tilde{p}$. Then

$$z_i(q) := \frac{q_i \tilde{p}_i}{p_i} \left(\sum_{k=1}^{n} \frac{q_k \tilde{p}_k}{p_k} \right)^{-1}.$$

Lastly, let $H : \Delta\Omega \to \mathbb{R}_+$ be strictly concave and define, for all $\tau \in T_p$,

$$C(\tau) := H(\tilde{p}) - \sum_q \tau(q) H\Big(z(q)\Big).$$

We can now assume, following Gentzkow and Kamenica (2014), that each $\tau \in T_p$ is associated with a cost $C(\tau) - C(\check{\tau})$. As R is concave, the more informative the splitting the greater the cost. We obtain, for $H(x) = -[x \log x + (1-x)\log(1-x)]$, a cost proportional to relative Shannon entropy.

Next, let

$$\Pi(q, w) := v(q) - \Big(\frac{mw}{mw-1}\Big) H\Big(z(q)\Big),$$

and

$$R_H(w) := \Big(\frac{mw}{mw-1}\Big)\left(\sum_q \check{\tau}(q) H\Big(z(q)\Big)\right).$$

The reduced problem of the principal then becomes

$$\max_{\{q^\dagger, \hat{q}\}} \quad w(q^\dagger, \hat{q})\Pi\Big(q^\dagger, w(q^\dagger, \hat{q})\Big) + \Big(1 - w(q^\dagger, \hat{q})\Big)\hat{\Pi}\Big(\hat{q}, w(q^\dagger, \hat{q})\Big) - R_H\Big(w(q^\dagger, \hat{q})\Big),$$

where $\hat{\Pi}(\cdot, w)$ denotes the concavification of the function $\Pi(\cdot, w)$.

7 Conclusion

This paper developed a general model of information design with agency. A principal chooses an information structure with a view to influencing a receiver, and hires an agent to implement said information structure. The effort of the agent cannot be contracted upon, and so the principal incentivizes the agent by rewarding her based on the action of the receiver. The principal then faces a trade-off between designing an information structure maximizing non-monetary payoffs, and minimizing the information rent that must be conceded to the agent in order to induce effort. While stylized, this general model captures for instance the broad features of the problem facing political organizations hiring public relations companies in order to sway public opinion.

Appendix

Proof of Lemma 1: Suppose $(\tilde{\tau}, \tilde{t}, \tilde{\alpha})$ solves (PP$_0$). Define $\tilde{f}(a_j) := \sum_{q:\tilde{\alpha}(q)=a_j} \tilde{\tau}(q)$, $\overline{\phi} := \max_{a_j \in A} \tilde{f}(a_j)$, and denote by $\dot{J}$ (respectively, $\ddot{J}$) the set of indices j such that $\tilde{f}(a_j) \neq 0$ (respectively, $\tilde{f}(a_j) = \overline{\phi}$). Then $(\tilde{\tau}, t, \tilde{\alpha})$ solves (PP$_0$) if and only if t solves

$$\min_t \sum_{j=1}^{m} \tilde{f}(a_j) t(a_j), \tag{CM}$$
$$\text{s.t.} \begin{cases} t(a_j) \geq 0, \ \forall a_j \in A, \\ \sum_{j=1}^{m} \tilde{f}(a_j) t(a_j) - c \geq \frac{1}{m} \sum_{j=1}^{m} t(a_j). \end{cases}$$

Next, note that t solves (CM) if and only if (i) $t(a_j) = 0$ for all $j \notin \dot{J}$, and (ii) $\left(\tilde{f}(a_1)t(a_1), \ldots, \tilde{f}(a_m)t(a_m)\right)$ solves

$$\min_{z \in \mathbb{R}^m} \sum_{j \in \dot{J}} z_j, \tag{CMZ}$$
$$\text{s.t.} \begin{cases} z_j \geq 0, \ \forall j, \\ z_j = 0 \text{ for all } j \notin \dot{J}, \\ \sum_{j \in \dot{J}} z_j \left(1 - \frac{1}{m\tilde{f}(a_j)}\right) \geq c. \end{cases}$$

However, z solves (CMZ) if and only if $z_j = 0$ for all $j \notin \ddot{J}$ and $\sum_{j \in \ddot{J}} z_j = cm\overline{\phi}/(m\overline{\phi} - 1)$.

Combining the previous remarks, we see that $(\tilde{\tau}, t, \tilde{\alpha})$ solves (PP$_0$) if and only if $t(a_j) = 0$ for all $j \notin \ddot{J}$ and $\sum_{j \in \ddot{J}} t(a_j) = cm/(m\overline{\phi} - 1)$. ■

Proof of Lemma 2: Pick $\Xi = (\tau, t, \alpha)$ solving (PP$_0$), and satisfying $t(a_j) = 0$, $\forall\ a_j \in A \setminus \{a^\dagger\}$, for some $a^\dagger \in A$ (we know such a plan exists, by Lemma 1). Moreover, suppose there exists $\dot{q} \in \operatorname{supp} \tau$ such that $\alpha(\dot{q}) \notin \overline{\overline{A}}(\dot{q})$ (otherwise there is nothing to prove). The following notation will be used throughout the proof: let $\overline{\overline{a}}$ denote a given action in $\overline{\overline{A}}(\dot{q})$, and $Q(\overline{\overline{a}}) := \{q \in \operatorname{supp} \tau \mid \alpha(q) = \overline{\overline{a}}\}$. Our goal will be to show that we can construct a plan solving (PP$_0$) and satisfying the condition stated in the lemma.

We first establish the following claim: there exists $\ddot{q} \in \operatorname{supp} \tau \setminus \{\dot{q}\}$ such that $\alpha(\ddot{q}) = \alpha(\dot{q})$. Suppose by way of contradiction that the claim is false. There are two cases to consider: $\alpha(\dot{q}) = a^\dagger$ (call this case 1), and $\alpha(\dot{q}) \neq a^\dagger$ (call this case 2). In case 1, consider $\tilde{\Xi} = (\tilde{\tau}, \tilde{t}, \tilde{\alpha})$ such that (i) $\tilde{\tau} = \tau$, (ii) $\tilde{\alpha}(q) = \alpha(q)$ for all $q \neq \dot{q}$, (iii) $\tilde{\alpha}(\dot{q}) = \overline{\overline{a}}$, (iv) $\tilde{t}\left(\overline{\overline{a}}(\dot{q})\right) = mc\Big[m\left(\tilde{\tau}(\dot{q}) +\right.$

$\sum_{q\in Q(\overline{\overline{a}})}\tilde{\tau}(q)\Big)-1\Big]^{-1}$, (v) $\tilde{t}(a_j)=0$ for all $a_j\neq\overline{\overline{a}}$. It is then easy to check that $\tilde{\Xi}$ is a feasible plan that yields to the principal a strictly greater expected payoff than Ξ, thus contradicting the optimality of the latter plan. Case 2, is further divided into two subcases, depending on whether $\overline{\overline{a}}=a^{\dagger}$ (call this subcase 2.1) or $\overline{\overline{a}}\neq a^{\dagger}$ (call this subcase 2.2). In subcase 2.1, consider $\tilde{\Xi}=(\tilde{\tau},\tilde{t},\tilde{\alpha})$ such that (i) $\tilde{\tau}=\tau$, (ii) $\tilde{\alpha}(q)=\alpha(q)$ for all $q\neq\dot{q}$, (iii) $\tilde{\alpha}(\dot{q})=\overline{\overline{a}}$, (iv) $\tilde{t}\Big(\overline{\overline{a}}(\dot{q})\Big)=mc\Big[m\Big(\tilde{\tau}(\dot{q})+\sum_{q\in Q(\overline{\overline{a}})}\tilde{\tau}(q)\Big)-1\Big]^{-1}$, (v) $\tilde{t}(a_j)=0$ for all $a_j\neq\overline{\overline{a}}$. It is again easy to check that $\tilde{\Xi}$ is a feasible plan that yields to the principal a strictly greater expected payoff than Ξ, thus contradicting the optimality of the latter plan. Lastly, in subcase 2.2, consider $\tilde{\Xi}=(\tilde{\tau},\tilde{t},\tilde{\alpha})$ such that (i) $\tilde{\tau}=\tau$, (ii) $\tilde{\alpha}(q)=\alpha(q)$ for all $q\neq\dot{q}$, (iii) $\tilde{\alpha}(\dot{q})=\overline{\overline{a}}$, (iv) $\tilde{t}=t$. The plan $\tilde{\Xi}$ is clearly feasible, and yields to the principal a strictly greater expected payoff than Ξ, contradicting once more the optimality of the latter plan. This finishes to show that the initial claim is true, namely, that there exists $\ddot{q}\in\operatorname{supp}\tau\setminus\{\dot{q}\}$ such that $\alpha(\ddot{q})=\alpha(\dot{q})$.

Next, consider the plan $\tilde{\Xi}=(\tilde{\tau},\tilde{t},\tilde{\alpha})$ obtained from Ξ by "merging" $\dot{q}$ and $\ddot{q}$ into

$$\dddot{q}:=\frac{\tau(\dot{q})\dot{q}+\tau(\ddot{q})\ddot{q}}{\tau(\dot{q})+\tau(\ddot{q})}$$

in the following way: (i) $\tilde{\tau}(\dddot{q})=\tau(\dot{q})+\tau(\ddot{q})+\tau(\dddot{q})$, (ii) $\tilde{\tau}(\dot{q})=\tilde{\tau}(\ddot{q})=0$, (iii) $\tilde{\tau}(q)=\tau(q)$ for all $q\notin\{\dot{q},\ddot{q},\dddot{q}\}$, (iv) $\tilde{\alpha}(\dddot{q})=\alpha(\dot{q})$, (v) $\tilde{\alpha}(q)=\alpha(q)$ for all $q\neq\dddot{q}$, (vi) $\tilde{t}=t$. The subset of posterior beliefs on which the action $\alpha(\dot{q})$ is receiver optimal is convex; hence, $\alpha(\dot{q})\in\overline{A}(\dddot{q})$. The plan $\tilde{\Xi}$ is thus feasible, and yields to the principal the same expected payoff as Ξ. So $\tilde{\Xi}$ solves (PP_0). If $\tilde{\alpha}(q)\in\overline{\overline{A}}(q)$ for all $q\in\operatorname{supp}\tilde{\tau}$, then we are done. Otherwise, repeating the previous steps with $\tilde{\Xi}$ instead of Ξ gives us a solution $\tilde{\Xi}'$ comprising a splitting $\tilde{\tau}'$ whose support contains one less element than $\tilde{\tau}$. As $\operatorname{supp}\tilde{\tau}$ is finite, continuing this way eventually yields a solution of (PP_0) satisfying the condition stated in the lemma. ■

Proof of Lemma 3: Pick a plan $\Xi=(\tau,t,\alpha)$ solving (PP_0) and satisfying the condition stated in Lemma 2. Repeating the steps in the proof of Lemma 1 shows that we may, without loss of generality, assume that there exists $a^{\dagger}\in A$ such that $t(a_j)=0$ for all $a_j\in A\setminus\{a^{\dagger}\}$. Lastly, as the subset of posterior beliefs on which the action $a^{\dagger}$ is receiver optimal is convex, we may, again without loss of generality, assume that there exists a unique $q^{\dagger}$ such that $\alpha(q^{\dagger})=a^{\dagger}$. For such a plan Ξ, the maximand in (PP_0) can be expressed as $\sum_q\tau(q)v(q)-\Big[R\Big(w(\tau)\Big)+c\Big]$.

Next, pick $\tilde{\tau}$ solving (PP_1), $\tilde{q}^{\dagger}\in\arg\max_q\tilde{\tau}(q)$ and $\tilde{a}^{\dagger}\in\overline{\overline{A}}(\tilde{q}^{\dagger})$. Now let $\tilde{\Xi}=(\tilde{\tau},\tilde{t},\tilde{\alpha})$ denote the plan comprising $\tilde{\tau}$ as well as transfers and recommendations such that (i) $\tilde{t}(\tilde{a}^{\dagger})=mc/(mw(\tilde{\tau})-1)$, (ii) $\tilde{t}(a_j)=0$ for all $a_j\neq\tilde{a}^{\dagger}$, (iii) $\tilde{\alpha}(q)\in\overline{\overline{A}}(q)$ for all $q\in\Delta\Omega$ and

$\tilde{\alpha}(\tilde{q}^\dagger) = \tilde{a}^\dagger$. This plan is clearly feasible, and yields to the principal an expected payoff equal to $\sum_q \tilde{\tau}(q)v(q) - \Big[R\Big(w(\tilde{\tau})\Big) + c\Big]$. Hence:

$$\begin{aligned}
\sum_q \tilde{\tau}(q)\Big[\sum_{i=1}^n q_i v\Big(\tilde{\alpha}(q),\omega_i\Big) - \tilde{t}\Big(\tilde{\alpha}(q)\Big)\Big] &= \sum_q \tilde{\tau}(q)v(q) - \Big[R\Big(w(\tilde{\tau})\Big) + c\Big] \\
&\geq \sum_q \tau(q)v(q) - \Big[R\Big(w(\tau)\Big) + c\Big] \\
&= \sum_q \tau(q)\Big[\sum_{i=1}^n q_i v\Big(\alpha(q),\omega_i\Big) - t\Big(\alpha(q)\Big)\Big].
\end{aligned}$$

Since Ξ solves (PP_0), and $\tilde{\Xi}$ is feasible, we conclude that $\tilde{\Xi}$ solves (PP_0) as well. ■

Lemma 8. *Consider $w > \frac{1}{m}$. If $\tau \in \arg\max_{\tau \in T_p} \Big\{\sum_q \tau(q)v(q) \;\Big|\; w(\tau) \geq w\Big\}$ then $M(w) = \hat{v}(p) - \sum_q \tau(q)v(q)$.*

Proof: Pick $\tau \in \arg\max_{\tau \in T_p} \Big\{\sum_q \tau(q)v(q) \;\Big|\; w(\tau) \geq w\Big\}$, and let α denote an arbitrary action recommendation function satisfying $\alpha(q) \in \overline{\overline{q}}$ for all $q \in \Delta\Omega$. Next, define $Q(a_j) := \{q \in \text{supp } \tau \mid \alpha(q) = a_j\}$, and let $a^\dagger \in \arg\max_{a_j \in A} \sum_{q \in Q(a_j)} \tau(q)$. Lastly, let $t(a^\dagger) = mc\Big[m\sum_{q \in Q(a^\dagger)} \tau(q) - 1\Big]^{-1}$ and $t(a_j) = 0$ for all $a_j \neq a^\dagger$. The plan $\Xi = (\tau, t, \alpha)$ is feasible, with $D(\Xi) = \hat{v}(p) - \sum_q \tau(q)v(q)$. Hence,

$$M(w) \leq \hat{v}(p) - \sum_q \tau(q)v(q). \tag{5}$$

Now pick

$$\tilde{\Xi} \in \arg\min \Big\{D(\Xi) \mid \; \Xi = (\tau, t, \alpha) \text{ is feasible and } w(\tau) \geq w\Big\},$$

that is, such that $D(\tilde{\Xi}) = M(w)$. Let $\tilde{Q}(a_j) := \{q \in \text{supp } \tilde{\tau} \mid \tilde{\alpha}(q) = a_j\}$, and define, for all j such that $\tilde{Q}(a_j) \neq \emptyset$,

$$q^j := \frac{\sum_{q \in \tilde{Q}(a_j)} \tilde{\tau}(q)q}{\sum_{q \in \tilde{Q}(a_j)} \tilde{\tau}(q)}.$$

Lastly, define $\dot{\tau} \in \Delta\Delta\Omega$ by $\dot{\tau}(q^j) = \sum_{q \in \tilde{Q}(a_j)} \tilde{\tau}(q)$. As the subset of posterior beliefs on which a given action a_j is receiver optimal is convex, notice that

$$D(\tilde{\Xi}) \geq \hat{v}(p) - \sum_q \dot{\tau}(q)v(q). \tag{6}$$

Combining (5) and (6) yields

$$\hat{v}(p) - \sum_q \dot{\tau}(q)v(q) \leq D(\tilde{\Xi}) = M(w) \leq \hat{v}(p) - \sum_q \tau(q)v(q).$$

Yet, as is easily checked: $\dot{\tau} \in T_p$ and $w(\dot{\tau}) \geq w$. Hence, by definition of τ: $\hat{v}(p) - \sum_q \tau(q)v(q) \leq \hat{v}(p) - \sum_q \dot{\tau}(q)v(q)$. It ensues that all weak inequalities in the last highlighted expression can be replaced by equalities. ■

Proof of Proposition 1: Let w solve (1). Then note that we must have $w > \frac{1}{m}$. Pick

$$\tau \in \arg\max_{\tau \in T_p} \left\{ \sum_q \tau(q)v(q) \,\middle|\, w(\tau) \geq w \right\}.$$

As R is a strictly decreasing function, applying Lemma 8 yields

$$M(w) + R(w) \geq \hat{v}(p) - \sum_q \tau(q)v(q) + R\Big(w(\tau)\Big) \tag{7}$$

We claim that τ solves (PP_1). Suppose by way of contradiction that the claim is false, and let $\tilde{\tau}$ do better than τ for the problem (PP_1):

$$\sum_q \tilde{\tau}(q)v(q) - R\Big(w(\tilde{\tau})\Big) > \sum_q \tau(q)v(q) - R\Big(w(\tau)\Big). \tag{8}$$

By Lemma 8:

$$M\Big(w(\tilde{\tau})\Big) + R\Big(w(\tilde{\tau})\Big) \leq \hat{v}(p) - \sum_q \tilde{\tau}(q)v(q) + R\Big(w(\tilde{\tau})\Big). \tag{9}$$

Then combining (7), (8) and (9) yields

$$\begin{aligned} M\Big(w(\tilde{\tau})\Big) + R\Big(w(\tilde{\tau})\Big) &\leq \hat{v}(p) - \sum_q \tilde{\tau}(q)v(q) + R\Big(w(\tilde{\tau})\Big) \\ &< \hat{v}(p) - \sum_q \tau(q)v(q) + R\Big(w(\tau)\Big) \\ &\leq M(w) + R(w), \end{aligned}$$

contradicting the optimality of w for (1). This establishes the claim that τ solves (PP_1). By Lemma 3, τ then also forms part of a plan Ξ solving (PP_0). ■

Proof of Proposition 2: Pick a plan $\Xi = (\tau, t, \alpha)$ solving (PP_0). Let $Q(a_j) := \{q \in \text{supp } \tau \mid \alpha(q) = a_j\}$, and define, for all j such that $Q(a_j) \neq \emptyset$,

$$q^j := \frac{\sum_{q \in Q(a_j)} \tau(q) q}{\sum_{q \in Q(a_j)} \tau(q)}.$$

Now consider $\tilde{\Xi} = (\tilde{\tau}, \tilde{t}, \tilde{\alpha})$ such that (i) $\tilde{\tau}(q^j) = \sum_{q \in Q(a_j)} \tau(q)$, (ii) $\tilde{\alpha}(q) \in \overline{\overline{A}}(q)$ for all $q \in \Delta\Omega$, (iii) $\tilde{t}(a^\dagger) = mc\Big[m \sum_{q:\tilde{\alpha}(q)=a^\dagger} \tilde{\tau}(q) - 1\Big]^{-1}$ for some action $a^\dagger \in \arg\max_{a_j \in A} \sum_{q:\tilde{\alpha}(q)=a_j} \tilde{\tau}(q)$, (iii) $\tilde{t}(a_j) = 0$ for all $a_j \neq a^\dagger$. Clearly, $\tilde{\Xi}$ is a feasible plan. Moreover, as the subset of posterior beliefs on which a given action a_j is receiver optimal is convex,

$$\sum_q \tilde{\tau}(q) v(q) = \sum_q \sum_{i=1}^n \tilde{\tau}(q) q_i v\Big(\tilde{\alpha}(q), \omega_i\Big) \geq \sum_q \sum_{i=1}^n \tau(q) q_i v\Big(\alpha(q), \omega_i\Big). \tag{10}$$

Lastly, since

$$\max_{a_j \in A} \mathbb{P}(a = a_j \mid \tilde{\Xi}) \geq \max_{a_j \in A} \mathbb{P}(a = a_j \mid \Xi), \tag{11}$$

similar arguments to those in the proof of Lemma 1 establish

$$\sum_q \tau(q) t\Big(\alpha(q)\Big) \geq \sum_q \tilde{\tau}(q) \tilde{t}\Big(\tilde{\alpha}(q)\Big). \tag{12}$$

As Ξ is optimal for (PP_0), the previous observations imply that the inequalities in (10), (11) and (12) must hold with equality. Hence (i) $D(\Xi) = D(\tilde{\Xi})$, and (ii) the principal's expected payoff under Ξ can be written as $\sum_q \tilde{\tau}(q) v(q) - \Big[R\Big(w(\tilde{\tau})\Big) + c\Big]$.

Next, let w solve (1), and pick $\dot{\tau} \in \arg\max_{\tau \in T_p} \Big\{\sum_q \tau(q) v(q) \;\Big|\; w(\tau) \geq w\Big\}$. Proposition 1, Lemma 8 and the remarks in the proof of Lemma 3 show that $\dot{\tau}$ forms part of a plan $\dot{\Xi}$ solving (PP_0) that yields to the principal an expected payoff that can be expressed as

$\sum_q \dot{\tau}(q)v(q) - \Big[R\Big(w(\dot{\tau})\Big) + c\Big] = \hat{v}(p) - M(w) - [R(w) + c]$. We then obtain

$$\begin{aligned}\sum_q \dot{\tau}(q)v(q) - \Big[R\Big(w(\dot{\tau})\Big) + c\Big] &= \hat{v}(p) - M(w) - [R(w) + c] \\ &\geq \hat{v}(p) - M\Big(w(\tilde{\tau})\Big) - \Big[R\Big(w(\tilde{\tau})\Big) + c\Big] \\ &\geq \hat{v}(p) - D(\tilde{\Xi}) - \Big[R\Big(w(\tilde{\tau})\Big) + c\Big] \\ &= \sum_q \tilde{\tau}(q)v(q) - \Big[R\Big(w(\tilde{\tau})\Big) + c\Big].\end{aligned}$$

Since the first and last expressions in the sequence above both equal the maximum value of (PP_0), all inequalities must hold with equality. As $D(\Xi) = D(\tilde{\Xi})$, we conclude that $D(\Xi) = L(\tilde{w})$ for some $\tilde{w}$ solving (1).

Now suppose $c_2 > c_1$, and that Ξ_i solves (PP_0) for $c = c_i$, $i = 1, 2$. Then, applying the previous remark gives $\tilde{w}_1$ and $\tilde{w}_2$ such that $\tilde{w}_i$ solves (1) for $c = c_i$, and $D(\Xi_i) = M(\tilde{w}_i)$. Moreover, the Monotone Selection Theorem of Milgrom and Shannon (1994) assures that $\tilde{w}_2 \geq \tilde{w}_1$. Since L is non-decreasing, we thus obtain

$$D(\Xi_2) = M(\tilde{w}_2) \geq M(\tilde{w}_1) = D(\Xi_1),$$

which shows that raising c results in a weakly greater distortion loss.

The second part of the proposition (raising c can result in a lower information rent for the agent) is shown by way of the following example with $n = 2$ states, and $m = 4$ actions. The prior belief $p = (3/4, 1/4)$. To simplify notation, in this example let q denote the posterior probability attached to $\omega = \omega_2$. The receiver-optimal actions mapping $\overline{A} : \Delta\Omega \to \mathcal{P}$ is given by $\overline{A}(0) = a_1$, $\overline{A}(1/4) = a_3$, $\overline{A}(1) = a_4$, and $\overline{A}(q) = a_2$ for all $q \notin \{0, 1/4, 1\}$. The principal's preferences are state-independent and summarized by $v(a_1) = v(a_4) = 1$, $v(a_2) = 0$ and $v(a_3) = 1 - \varepsilon$, where $0 < \varepsilon < 1/4$. Then, as is easily checked, either $c < 6\varepsilon$ in which case the unique solution of (PP_0) splits the prior on 0 and 1 and pays the agent if and only if $a = a_1$, or $c > 6\varepsilon$ in which case the unique solution generates no information and pays the agent if and only if $a = a_3$. Consequently, the agent's information rent is 1/2 for $c < 6\varepsilon$ but only 1/3 for $c > 6\varepsilon$. ■

Proof of Lemma 4: Follows from the Monotone Selection Theorem of Milgrom and Shannon (1994). ■

Proof of Proposition 3: As $c < \bar{c}$, $\mathbf{1}_p$ does not form part of a plan solving (PP$_0$). Hence, by Lemma 3, $\mathbf{1}_p$ does not solve (PP$_1$). Let $\{q^\dagger, \hat{q}\}$ denote a solution of (PP$_2$), $\hat{\tau} \in T_{\hat{q}}^{\hat{v}}$, $a^\dagger \in \overline{\overline{A}}(q^\dagger)$ and $\tau = w(q^\dagger, \hat{q})\mathbf{1}_{q^\dagger} + \Big(1 - w(q^\dagger, \hat{q})\Big)\hat{\tau}$. Then

$$\sum_q \tau(q)v(q) - R\Big(w(\tau)\Big) \geq w(q^\dagger, \hat{q})v(q^\dagger) + \Big(1 - w(q^\dagger, \hat{q})\Big)\hat{v}(\hat{q}) - R\Big(w(q^\dagger, \hat{q})\Big). \tag{13}$$

We claim that τ solves (PP$_1$). Suppose by way of contradiction that the claim is false, and let $\tilde{\tau} \in T_p$ do better than τ for the problem (PP$_1$):

$$\sum_q \tilde{\tau}(q)v(q) - R\Big(w(\tilde{\tau})\Big) > \sum_q \tau(q)v(q) - R\Big(w(\tau)\Big). \tag{14}$$

Pick $\tilde{q}^\dagger \in \arg\max_q \tilde{\tau}(q)$, define $\tilde{w} := \tilde{\tau}(\tilde{q}^\dagger)$, and let $\hat{\tilde{q}} := \frac{p - \tilde{w}\tilde{q}^\dagger}{1 - \tilde{w}}$. Then $\{\tilde{q}^\dagger, \hat{\tilde{q}}\}$ is a binary splitting of p, $\tilde{w} = w(\tilde{q}^\dagger, \hat{\tilde{q}})$ and, by writing

$$\sum_q \tilde{\tau}(q)v(q) = \tilde{w}v(\tilde{q}^\dagger) + (1 - \tilde{w}) \sum_{q \neq \tilde{q}^\dagger} \frac{\tilde{\tau}(q)}{1 - \tilde{w}} v(q),$$

we see that

$$\begin{aligned}\sum_q \tilde{\tau}(q)v(q) - R\Big(w(\tilde{\tau})\Big) &\leq \tilde{w}v(\tilde{q}^\dagger) + (1 - \tilde{w})\hat{v}(\hat{\tilde{q}}) - R(\tilde{w}) \\ &= w(\tilde{q}^\dagger, \hat{\tilde{q}})v(\tilde{q}^\dagger) + \Big(1 - w(\tilde{q}^\dagger, \hat{\tilde{q}})\Big)\hat{v}(\hat{\tilde{q}}) - R\Big(w(\tilde{q}^\dagger, \hat{\tilde{q}})\Big). \end{aligned} \tag{15}$$

Combining (13), (14) and (15) then gives

$$w(\tilde{q}^\dagger, \hat{\tilde{q}})v(\tilde{q}^\dagger) + \Big(1 - w(\tilde{q}^\dagger, \hat{\tilde{q}})\Big)\hat{v}(\hat{\tilde{q}}) - R\Big(w(\tilde{q}^\dagger, \hat{\tilde{q}})\Big) > w(q^\dagger, \hat{q})v(q^\dagger) + \Big(1 - w(q^\dagger, \hat{q})\Big)\hat{v}(\hat{q}) - R\Big(w(q^\dagger, \hat{q})\Big),$$

contradicting the optimality of $\{q^\dagger, \hat{q}\}$ for the problem (PP$_2$). This establishes the claim that τ solves (PP$_1$). Repeating the arguments used in the proof of Lemma 3 then shows that the plan $\Xi = (\tau, t, \alpha)$ solves (PP$_0$), where transfers and recommendations are determined by (i) $\alpha(q) \in \overline{\overline{A}}(q)$ for all $q \in \Delta\Omega$ and $\alpha(q^\dagger) = a^\dagger$, (ii) $t(a^\dagger) = mc\Big[m\sum_{q:\alpha(q)=a^\dagger} \tau(q) - 1\Big]^{-1}$, (iii) $t(a_j) = 0$ for all $a_j \neq a^\dagger$. ■

Proof of Proposition 4: Consider $c < \bar{c}$. By Lemma 5, any solution $\{q^\dagger, \hat{q}\}$ of the reduced problem (PP$_2$) is such that $\hat{q} = 0$ and $q^\dagger = \frac{1}{2} + x$, for some $x \in (0, 1/2)$. Using $w(\frac{1}{2} + x, 0) =$

$\frac{1}{2x+1} =: w$, for such a binary splitting the maximand in (PP$_2$) can be written as

$$3w(\tfrac{1}{2}+x)+(1-w)(\tfrac{1}{2}+x)-\tfrac{c}{2w-1}=\tfrac{5}{2}-w-\tfrac{c}{2w-1}.$$

The first-order condition on w gives

$$\frac{2c}{(2w-1)^2}-1=0.$$

Thus $w=\frac{1}{2}+\sqrt{\frac{c}{2}}$ and $x=\frac{1-\sqrt{2c}}{2(1+\sqrt{2c})}$. ■

Proof of Proposition 5: If $c>\bar{c}$ then $\tau^*=\mathbf{1}_p$, so consider $c<\bar{c}$ in the rest of the proof. Applying Lemmata 6 and 7 and the observations made in Section 4, we may retrieve τ^* by direct comparison of (a) the principal's expected payoff resulting from the optimal splitting $\{q^\dagger,\hat{q}\}$ for problem (PP$_2$) within the class of binary splittings satisfying $q^\dagger=\overline{q}$ and $\hat{q}\in[0,1-\overline{q}]$, and (b) the principal's expected payoff resulting from the optimal splitting $\{q^\dagger,\hat{q}\}$ within the class of binary splittings satisfying $q^\dagger>1-\overline{q}$ and $\hat{q}=1$. Below, we first find (a), then (b).

The optimal splitting in (a) is entirely determined by $w(q^\dagger,\hat{q})=w(\overline{q},\hat{q})=:w$, via $\hat{q}=\frac{1-2w\overline{q}}{2(1-w)}$. Note that $\hat{q}\in[0,1-\overline{q}]\Leftrightarrow w\in[\frac{1}{2},\frac{1}{2\overline{q}}]$. Therefore, the optimal splitting in (a) solves:

$$\max_{w\in[\frac{1}{2},\frac{1}{2\overline{q}}]} \quad w+\left(\frac{1-2w\overline{q}}{2(1-\overline{q})}\right)v_M-\frac{c}{3w-1}.$$

The maximand above is a concave function of w. The first-order condition yields

$$w(c)=\frac{1}{3}\left(\sqrt{\frac{3c}{\frac{\overline{q}}{1-\overline{q}}v_M-1}}+1\right).$$

Define c_1 and c_2 by $w(c_1)=1/2$ and $w(c_2)=1/(2\overline{q})$; then $c_2>c_1$, since w is increasing. Straightforward algebra gives $c_1=\frac{1}{12}\left(\frac{\overline{q}}{1-\overline{q}}v_M-1\right)$ and $c_2=\frac{1}{3}\left(\frac{3}{2\overline{q}}-1\right)^2\left(\frac{\overline{q}}{1-\overline{q}}v_M-1\right)$.

Similarly, the optimal splitting in (b) is entirely determined by $w(q^\dagger,\hat{q})=w(q^\dagger,1)=:W$, via $q^\dagger=\frac{2W-1}{2W}$. Note that $q^\dagger\in(1-\overline{q},1/2)\Leftrightarrow W\in\left(\frac{1}{2\overline{q}},1\right)$. Therefore, the optimal splitting in (b) solves

$$\max_{W\in\left(\frac{1}{2\overline{q}},1\right)} \quad Wv_M+(1-W)-\frac{c}{3W-1}.$$

The first-order condition yields

$$W(c) = \frac{1}{3}\left(\sqrt{\frac{3c}{1-v_M}} + 1\right).$$

Lastly, define $\tilde{c}$ by $W(\tilde{c}) = 1/(2\overline{q})$. Then $\tilde{c} = \frac{1}{3}\left(\frac{3}{2\overline{q}-1}\right)(1-v_M)$. One checks that $\tilde{c} > c_2 \Leftrightarrow 2(1-\overline{q}) > v_M$. This shows $\tilde{c} > c_2$. This remark, in turn, establishes that, if $\{q^\dagger, 1\}$ solves (PP_2) and satisfies $q^\dagger > 1-\overline{q}$, then $c > c_2$. So the cutoff c_3 in the statement of the proposition must be strictly greater than c_2. ■

References

Angelucci, C. (2017) Motivating agents to acquire information, *mimeo*.

Aumann, R. J., Maschler, M. and Stearns, R. E. (1995) *Repeated games with incomplete information*, MIT press.

Bizzotto, J. and Vigier, A. (2021) School design, *mimeo*.

Boleslavsky, R. and Kim, K. (2020) Bayesian persuasion and moral hazard, *mimeo*.

Boleslavsky, R., Shadmehr, M. and Sonin, K. (2020) Media Freedom in the Shadow of a Coup, *Journal of the European Economic Association*, jvaa040.

Carroll, G. (2019) Robust incentives for information acquisition, *Journal of Economic Theory*.

Gentzkow, M. and Kamenica, E. (2014) Costly persuasion, *American Economic Review*, **104**, 457–462.

Gibson, R. and Rommele, A. (2009) Measuring the Professionalization of Political Campaigns, *Party Politics*, **15**, 265–293.

Gratton, G. and Lee, B. E. (2020) Liberty, security, and accountability: The rise and fall of illiberal democracies.

Grusell, M. and Nord, L. (2020) Setting the trend or changing the game? professionalization and digitalization of election campaigns in sweden, *Journal of Political Marketing*, **19**, 258–278.

Herresthal, C. (2020) Hidden testing and selective disclosure of evidence, *mimeo*.

Hillygus, S. and Shields, T. (2009) *The Persuadable Voter : Wedge Issues in Presidential Campaigns*, Princeton University Press.

Kamenica, E. and Gentzkow, M. (2011) Bayesian persuasion, *American Economic Review*, **101**, 2590–2615.

Kleinberg, J. and Raghavan, M. (2020) How do classifiers induce agents to invest effort strategically?, *ACM Transactions on Economics and Computation (TEAC)*, **8**, 1–23.

Kolotilin, A., Mylovanov, T. and Zapechelnyuk, A. (2021) Censorship as optimal persuasion, *Available at SSRN 3501474*.

Kremer, I., Mansour, Y. and Perry, M. (2014) Implementing the "wisdom of the crowd", *Journal of Political Economy*, **122**, 988–1012.

Milgrom, P. and Shannon, C. (1994) Monotone comparative statics, *Econometrica*, **62**, 157–180.

Perez-Richet, E. and Skreta, V. (2020) Test design under falsification, *mimeo*.

Rappoport, D. and Somma, V. (2017) Incentivizing information design, *mimeo*.

Rodina, D. and Farragut, J. (2017) Inducing effort through grades, *mimeo*.

Rosar, F. (2017) Test design under voluntary participation, *Games and Economic Behavior*, **104**, 632–655.

Shadmehr, M. and Bernhardt, D. (2015) State censorship, *American Economic Journal: Microeconomics*, **7**, 280–307.

Sheingate, A. (2016) *Building a Business of Politics: The Rise of Political Consulting and the Transformation of American Democracy*, Oxford University Press.

Strömbäck, J. and Kiousis, S. (2012) Strategic Political Communication in Election Campaigns, *In Carsten Reinemann (Editor) Communication (109-128)*.

Szalay, D. (2005) The economics of clear advice and extreme options, *The Review of Economic Studies*, **72**, 1173–1198.

Tenscher, J., Koc-Michalska, K., Lilleker, D., Mykkänen, J., Walter, A., Findor, A. and Jalali, J., Carlos Róka (2016) The professionals speak: Practitioners' perspectives on professional election campaigning, *European Journal of Communication*, **31**, 95–119.

Yoder, N. (2020) Designing incentives for heterogeneous researchers, *mimeo*.

Zapechelnyuk, A. (2020) Optimal quality certification, *American Economic Review: Insights*, **2**, 161–176.

Zermeño, L. (2011) A principal-expert model and the value of menus, *mimeo*.