SEQUENTIAL OBFUSCATION AND TOXIC ARGUMENTATION

Andriy Zapechelnyuk

Abstract. We study an information design problem where two competing senders move sequentially. The second mover can either reveal more information (constructive argumentation) or obfuscate the first mover's information (toxic argumentation). We show that sequential obfuscation of an initially disclosed state never reveals more, and sometimes reveals strictly less information than sequential disclosure of an initially hidden state. Sequential obfuscation is completely uninformative when the senders are risk averse or risk neutral, or when they have zero-sum preferences. By contrast, sequential disclosure generally reveals some information, and it is fully revealing when the senders are risk neutral or have zero-sum preferences.

JEL Classification: D82, D83

Keywords: Information design, Bayesian persuasion, information structure, disclosure, obfuscation, garbling

---

Date: June 19, 2022.

Zapechelnyuk: School of Economics and Finance, University of St Andrews, Castlecliffe, the Scores, St Andrews KY16 9AR, UK. *E-mail:* az48@st-andrews.ac.uk.

The author thanks Peter Norman and Ina Taneva for very helpful comments.

1. Introduction

Toxic arguments are those that carry negative criticism, blame, and contempt. Negative criticism is deployed to point out that the opponent's arguments are inaccurate or unsound. Blame and contempt are used to target the opponent's personal flaws and to show that the opponent is not worth listening to. The ultimate purpose of toxic argumentation is to cast a doubt on the validity and credibility of the opponent's arguments and reduce their informational content as perceived by the audience. Toxic argumentation stands in contrast with constructive argumentation that adds to rather than subtracts from the informational content of the opponent's arguments.

In this paper we use the concept of information obfuscation, or garbling, to model toxic argumentation. In our model, two debaters sequentially choose information disclosure strategies of an uncertain state of the world in order to influence the choice of a listener. We compare two cases: *sequential disclosure* and *sequential obfuscation*. In the case of sequential disclosure, the second mover uses constructive arguments to reveal additional information about the state. In the case of sequential obfuscation, the second mover uses toxic arguments to obfuscate the information revealed by the first mover.

We ask how and to what extent the nature of counterarguments by the second mover affects information disclosure to the listener. At a glance, complementing one's argument with another informative argument should result in more information disclosure than obfuscating one's argument. But after a moment of reflection this should not be obvious. The first mover can adjust her behavior in anticipation of the opponent's counteraction. For example, she can strategically choose to disclose more information when expecting the opponent to obfuscate some of it. Furthermore, note that the case of sequential obfuscation can be equivalently represented literally, as two candidates sequentially obfuscating an initially revealed state of the world. So, there is an intrinsic symmetry between sequential disclosure of an initially hidden state and sequential obfuscation of an initially revealed state. There is no difference when there is only one sender, and it is not obvious what difference it makes to the strategic interaction of two senders.

We begin by showing how the problems of sequential disclosure and sequential obfuscation can be simplified. After the simplification, the difference between the two problems becomes apparent. In both cases, the first mover solves a constrained optimization problem with the same objective but different constraints. In sequential disclosure, the first mover chooses among the outcomes that the second mover cannot improve upon by further disclosure. In contrast, in sequential obfuscation, the first mover chooses among the outcomes that the second mover cannot improve upon by obfuscation. This allows us to show that sequential obfuscation cannot make the listener better informed than sequential disclosure. Moreover, we show that sequential obfuscation is completely uninformative when the debaters are weakly risk averse, or when they have zero-sum preferences. By contrast, sequential disclosure generally reveals some information, and it is fully reveals the state when the debaters are risk neutral, or when they have zero-sum preferences. Thus, in the special cases of risk neutral or zero-sum preferences of the debaters, the two problems yield the opposite extreme solutions.

Example. For illustration, consider two political candidates who compete in an election. They engage in a debate on an issue that decides the vote of the median voter. Let the information about the debated issue be summarized by a state $\theta \in [0, 1]$. The candidates sequentially appoint experts who reveal information about the state to the public. An expert is represented by an information structure that either reveals information about hidden state or obfuscates what has already been revealed. The candidates have access to a large pool of experts. All information structures are available to both candidates in sequential disclosure, but in sequential obfuscation candidate 2 is restricted to obfuscation.

Let the position of the median voter on the political spectrum be uncertain, so that the median voter votes for candidate 1 if and only if his position is below the expected state conditional on revealed information. Let $G(x)$ be the probability that the median voter votes for candidate 1 if the expected state is x, so G is interpreted as a cumulative distribution of the position of the median voter. Each candidate wishes to maximize their own probability of winning, so the utility of candidate 1 is $G(x)$, and the utility of candidate 2 is $1 - G(x)$.

In this example, irrespective of the distribution G of the position of the median voter, sequential obfuscation is completely uninformative, whereas sequential disclosure fully reveals the state. There is a simple intuition for this. Consider the case of sequential obfuscation. Because this is a constant-sum game between the two candidates, no outcome is Pareto superior (for the candidates) to the uninformative outcome. But because the uninformative outcome is enforceable individually by each candidate (by full obfuscation), this becomes the unique equilibrium outcome of sequential obfuscation. The argument that sequential disclosure fully reveals the state is analogous.

Related Literature. This paper is related to and contributes into the literature on competition in information design where senders commit to information disclosure protocols before learning the state of the world. Gentzkow and Kamenica (2017a,b), Li and Norman (2018), and Ravindran and Cui (2020) consider senders who simultaneously choose information structures. The peculiarity of simultaneous disclosure is that when more than one sender discloses the same bit of information, no sender can unilaterally prevent its disclosure. This leads to multiplicity of equilibria, in particular, full disclosure of the state is always an equilibrium. Mylovanov and Zapechelnyuk (2022) introduce an equilibrium refinement to obtain a unique equilibrium outcome of simultaneous disclosure. Boleslavsky and Cotton (2018) and Au and Kawai (2020) restrict the senders to disclose different coordinates of a multidimensional state, thus preventing the overlap in the information disclosure. Li and Norman (2021) consider sequential, rather than simultaneous disclosure, where sequential moves lead a unique equilibrium outcome.

As our paper compares sequential disclosure and sequential obfuscation, Li and Norman's (2021) study of sequential disclosure is the closest paper to ours. The setting of Li and Norman (2021) is more general than ours, and their focus is on the characterization of equilibria of sequential disclosure, and on the comparison with simultaneous disclosure and with disclosure by a single sender. Our result that the sequential obfuscation problem reduces to the first sender's constrained optimization (Proposition 1) is a direct adaptation of the correspondent result in Li and Norman (2021) that applies to sequential disclosure. The novelty of our paper is that we address sequential obfuscation and compare it with sequential disclosure. The additional structure relative to Li and Norman (2021) also allows us to obtain a new result when sequential disclosure fully reveals the state (Proposition 3).

To our knowledge, this paper is first to study information obfuscation in the role where it is distinct from information disclosure. When there is a single sender, obfuscation of an initially revealed state is strategically identical to disclosure of an initially hidden state. In the information design literature with a single sender, the term *obfuscation* (garbling, confusion) appears synonymously to information disclosure but is often used to emphasize the interpretation where the sender reduces information about an initially revealed state (e.g., Chan, Gupta, Li, and Wang, 2019; Edmond and Lu, 2021; Li, Song, and Zhao, 2022).

We adopt a so-called linear information design approach to modeling obfuscation. Linearity refers to the property that the payoffs depend on the posterior belief about the state only through the posterior mean. This approach received a lot of attention on the literature (Kamenica and Gentzkow, 2011; Gentzkow and Kamenica, 2016; Kolotilin, Mylovanov, Zapechelnyuk, and Li, 2017; Kolotilin, 2018; Kolotilin and Zapechelnyuk, 2019; Dworczak and Martini, 2019; Arieli, Babichenko, Smorodinsky, and Yamashita, 2022; Kleiner, Moldovanu, and Strack, 2021). It has been used in many applications of information design, including media control (Gehlbach and Sonin, 2014; Ginzburg, 2019; Gitmez and Molavi, 2020; Kolotilin, Mylovanov, and Zapechelnyuk, 2022), clinical trials (Kolotilin, 2015), voter persuasion (Alonso and Câmara, 2016), transparency benchmarks (Duffie, Dworczak, and Zhu, 2017), stress tests (Goldstein and Leitner, 2018; Orlov, Zryumov, and Skrzypach, 2020), online markets (Romanyuk and Smolin, 2019), attention management (Lipnowski, Mathevet, and Wei, 2020; Bloedel and Segal, 2020), and quality certification (Zapechelnyuk, 2020).

The fundamental assumption in information design, which is also adopted in this paper, is that the senders can commit to information structures ex ante, before learning any information about the state of the world. While this assumption is certainly restrictive, to a certain extent it is justified by Zapechelnyuk (2022) who shows the equivalence of implementable outcomes in the settings where the sender, whose preferences are monotone, is uninformed about the state and where she is informed about the state prior to committing to an information structure.[1]

Our paper is also related to the literature on informational lobbying, where a policy maker or legislator consults two or more biased experts. A focal question in this literature is whether seeking advice of more experts can improve the information of the policy maker. In Gilligan and Krehbiel (1989), Krishna and Morgan (2001a,b), Battaglini (2002), Ambrus and Takahashi (2008), Li (2010), and Mylovanov and Zapechelnyuk (2013a,b) the experts know the state of the world, so consulting multiple experts has no informational benefit, but it can improve the incentives for information disclosure. In Austen-Smith (1993), Wolinsky (2002), Battaglini (2004), Levy and Razin (2007), and Ambrus and Lu (2014), each expert's private information is partial, and consulting more that one expert can improve the informational content, whereas Li (2010) shows that more experts can result in less disclosure. The effects of the order in which experts present their arguments are explored in Krishna and Morgan (2001b) and D'Agostino and Seidmann (2022), and the collusion of the experts is explored in Zapechelnyuk (2013). Our paper contributes to this literature by addressing the complementary question about the effect of adding an "expert" who obfuscates existing information instead of enriching it.

[1] Other papers that study information design with privately informed sender include Perez-Richet (2014), Degan and Li (2016), Hedlund (2017), and Koessler and Skreta (2021).

2. Model

2.1. **Basic Setting.** There are a receiver and two senders. The receiver chooses an action in an interval $A = [a_0, a_1]$. Each sender $i = 1, 2$ obtains utility $u_i : A \to \mathbb{R}$ that depends only on the receiver's action. A state of the world θ is a real-valued random variable with a common prior μ_0. Assume that μ_0 has compact support whose lowest and highest points are 0 and 1, so $\theta \in \Theta = [0, 1]$. Given an expected value of the state $x \in \Theta$, the receiver is assumed to choose an action $a(x) \in A$, where $a(x)$ is weakly increasing.[2]

We assume that the sender's preferences are opposing and monotone in the receiver's action,

$$u_1'(a) > 0 \quad \text{and} \quad u_2'(a) < 0, \quad a \in A. \tag{1}$$

This assumption introduced for the convenience of interpretation of the results. It is formally required only for Proposition 4 in Section 3.3.

Let us describe the senders' strategies. Let M_i be a set of messages of sender $i = 1, 2$. Suppose that the sets M_1 and M_2 are rich enough, so $\Theta \subseteq M_1$ and $\Theta \times M_1 \subseteq M_2$. A strategy of sender 1 is a mapping $\phi_1 : \Theta \to \Delta(M_1)$ that associates with each state θ a conditional probability distribution $\phi_1(\cdot|\theta)$ over sender 1's messages in M_1. A strategy of sender 2 is a mapping $\phi_2 : \Theta \times M_1 \to \Delta(M_2)$ that associates with each state θ and each message m_1 of sender 1 a conditional probability distribution $\phi_2(\cdot|\theta, m_1)$ over sender 2's messages in M_2.

The timing is as follows. Senders 1 and 2 choose their strategies sequentially. Then state θ realizes. Then, message m_1 is generated according to sender 1's strategy, after which message m_2 is generated according to sender 2's strategy. The receiver observes the strategies of the senders and message m_2 of sender 2 (but not message m_1 of sender 1). Given the prior μ_0 and the observed information, the receiver derives the posterior expected state x, and chooses action $a(x)$.

Because the senders' utilities depend only on the the receiver's action, which in turn depends only on the expected state, the information disclosed by a message can be summarized by the probability distribution over the posterior expected state induced

[2]To interpret $a(x)$, consider a population of heterogeneous receivers who choose to support sender 1 or 2 depending on their belief about the expected state and their private type. Let $a(x)$ be the fraction of the population who support sender 1 when the expected state is x (and $1 - a(x)$ is the fraction of the population who support sender 2). So $a(x)$ captures the heterogeneity of the predisposition towards sender 1 in the population.

by this message. Given a pair of strategies (ϕ_1, ϕ_2), let $\mu_1(\phi_1) \in \Delta(\Theta)$ be the distribution of the expected state induced by messages of sender 1, and let $\mu_2(\phi_1, \phi_2) \in \Delta(\Theta)$ be the distribution of the expected state induced by messages of sender 2.

We compare distributions of the expected state by their Blackwell informativeness (Blackwell, 1953) for the receiver. We say that distribution μ' is *more informative* than distribution μ'', denoted by $\mu' \succeq \mu''$, if μ' is a mean preserving spread of μ''.

2.2. **Sequential Disclosure and Sequential Obfuscation.** We consider two variants of the basic setting: a model of sequential disclosure and a model of sequential obfuscation. These models impose different constraints on the strategy of sender 2.

In sequential disclosure, sender 2 reveals information in addition to what has been revealed by sender 1's message. That is, the receiver can always deduce m_1 from m_2. This formalism captures the idea that the receiver observes both messages, so sender 2 cannot hide what has been revealed by sender 1. By Blackwell (1953), this means that, given the distribution $\mu_1(\phi_1)$ of the expected state induced by sender 1's strategy ϕ_1, strategy ϕ_2 must induce a weakly more informative distribution, so ϕ_2 must satisfy

$$\mu_2(\phi_1, \phi_2) \succeq \mu_1(\phi_1).$$

In sequential obfuscation, sender 2 obfuscates (or garbles) information revealed by sender 1's message. That is, if the receiver was able to observe m_1 instead of m_2, he could deduce m_2. This means that, given the distribution $\mu_1(\phi_1)$ of the expected state induced by sender 1's strategy ϕ_1, strategy ϕ_2 must induce a weakly less informative distribution, so ϕ_2 must satisfy

$$\mu_1(\phi_1) \succeq \mu_2(\phi_1, \phi_2).$$

We are interested in the characterisation and comparison of equilibria in the models of sequential disclosure and sequential obfuscation. The solution concept is subgame prefect equilibrium.

3. Results

3.1. **Equilibrium outcomes.** An *outcome* μ of sequential disclosure or sequential obfuscation with a given pair of strategies (ϕ_1, ϕ_2) is the distribution of the posterior expected state induced by the message of sender 2, $\mu = \mu_2(\phi_1, \phi_2)$. The outcome summarizes the information revealed to the receiver. It also determines the expected utilities of the senders. Let $V_i(\mu)$ be the expected utility of sender i when the outcome is $\mu \in \Delta(\Theta)$,

$$V_i(\mu) = \int_{x \in \Theta} u_i(a(x)) \mathrm{d}\mu(x), \quad i = 1, 2.$$

Given a prior μ_0, an outcome $\mu \in \Delta(\Theta)$ is implementable by information structures, in particular, by sequential disclosure or sequential obfuscation, if and only if μ_0 is more informative than μ (Blackwell, 1953). Let $\mathcal{M}$ be the set of implementable outcomes,

$$\mathcal{M} = \{\mu \in \Delta(\Theta) : \mu_0 \succeq \mu\}.$$

We use the notion of unimprovable outcomes[3] to simplify the problems of finding subgame perfect equilibria in sequential disclosure and sequential obfuscation.

An implementable outcome $\mu \in \mathcal{M}$ is *unimprovable by disclosure* for sender i if she cannot be better off with any outcome μ' that can be obtained from μ by disclosure,

$$V_i(\mu) \geq V_i(\mu') \text{ for all } \mu' \in \mathcal{M} \text{ such that } \mu_0 \succeq \mu' \succeq \mu.$$

An outcome $\mu \in \mathcal{M}$ is *unimprovable by obfuscation* for sender i if she cannot be better off with any outcome μ' that can be obtained from μ by obfuscation,

$$V_i(\mu) \geq V_i(\mu') \text{ for all } \mu' \in \mathcal{M} \text{ such that } \mu \succeq \mu'.$$

Let $\mathcal{M}_2^D$ and $\mathcal{M}_2^O$ be the set of implementable outcomes that are unimprovable by disclosure and obfuscation, respectively, for sender 2.

We now show that the problem of sequential disclosure (sequential obfuscation) is equivalent to the problem where only sender 1 chooses an information structure. Because sender 2 is able to distort some choices of sender 1 by revealing (obfuscating) information, sender 1 can only attain outcomes that sender 2 does not want to improve upon. Sender 1 then chooses the best among such outcomes.

Consider two problems where sender 1 chooses an outcome to maximize her expected payoff among the outcomes that are unimprovable by disclosure and obfuscation, respectively, for sender 2:

$$\max_{\mu \in \mathcal{M}_2^D} V_1(\mu), \tag{P$_\mathrm{D}$}$$

$$\max_{\mu \in \mathcal{M}_2^O} V_1(\mu). \tag{P$_\mathrm{O}$}$$

[3]Variants of this notion appear in Gentzkow and Kamenica (2017b) and Li and Norman (2021).

Proposition 1. *An outcome $\mu \in \Delta(\Theta)$ is an equilibrium outcome of sequential disclosure (sequential obfuscation) if and only if it is a solution of problem* ($\mathrm{P_D}$) *(respectively,* ($\mathrm{P_O}$)*).*

Li and Norman (2021) prove the statement of Proposition 1 for sequential disclosure. The argument for sequential obfuscation is analogous. The idea behind Proposition 1 is reminiscent of the revelation principle. If an equilibrium of sequential obfuscation by two senders leads to an outcome μ, then it must remain equilibrium if sender 1 implements μ directly. Sender 2 then has no incentive to obfuscate μ, because if she did, she would have done so in the original equilibrium.

3.2. **Comparison of disclosure and obfuscation.** Proposition 1 illuminates the difference between disclosure and obfuscation. Loosely speaking, sequential disclosure restricts sender 1's choice to outcomes that are sufficiently revealing from sender 2's perspective, so that sender 2 does not wish to reveal any more. Similarly, sequential obfuscation restricts sender 1's choice to outcomes that are sufficiently unrevealing from sender 2's perspective, so that sender 2 does not wish to obfuscate them. The set of outcomes that are unimprovable by both disclosure and obfuscation for sender 2 has measure zero set for a generic decision function a. Thus, sender 1 optimizes on two essentially disjoint sets in the two problems, one clearly favoring more information disclosure than the other.

Let us now support the above argument by a formal result. It demonstrates that sequential obfuscation cannot be more informative than sequential disclosure.

Proposition 2. *Let μ^D and μ^O be equilibrium outcomes of sequential disclosure and sequential obfuscation, respectively, and suppose that the senders' expected utilities are not identical,*

$$\big(V_1(\mu^D), V_2(\mu^D)\big) \neq \big(V_1(\mu^O), V_2(\mu^O)\big) .$$

Then μ^O cannot be more informative than μ^D.

The proof is in Appendix A.1.

Next, we show that sequential disclosure typically reveals some information, except when both senders unanimously prefer to reveal none. Similarly, sequential obfuscation typically obfuscates some information, except when both senders unanimously prefer to fully reveal the state.

An outcome μ is called *no disclosure* if it reveals no information about the state, that is, it puts probability one on the prior expected value of the state.

An outcome μ is called *full disclosure* if it reveals the state, that is, $\mu = \mu_0$.

Corollary 1. *No disclosure (full disclosure) is an equilibrium outcome of sequential disclosure (sequential obfuscation, respectively) if and only if it is preferred to all other outcomes by both senders.*

Corollary 1 follows from Proposition 1 and the fact that the outcome of no disclosure is unimprovable by disclosure for sender 2 if only if it is sender 2's preferred outcome, and, similarly, the outcome of full disclosure is unimprovable by obfuscation for sender 2 if only if it is sender 2's preferred outcome.

Next, look at the opposite extreme. We provide a condition on the senders' utilities such that sequential disclosure fully reveals the state and sequential obfuscation reveals no information at all.

Suppose that the senders' utilities are linear functions of each other, so

$$u_2(y) = b - cu_1(y) \quad \text{for some } b \in \mathbb{R} \text{ and } c > 0. \tag{2}$$

This assumption generalizes two special cases that are prominent in the literature. It holds when the senders have zero-sum or constant-sum utilities. It also holds when the senders' utilities are linear functions of y, so the senders are risk neutral.

We show that under this assumption, sequential disclosure reveals the state and sequential obfuscation reveals no information.

Proposition 3. *Suppose that the senders' utilities satisfy (2). Then full disclosure (no disclosure) is an equilibrium outcome of sequential disclosure (sequential obfuscation, respectively). Moreover, this is the unique equilibrium outcome for a generic decision function $a(\cdot)$ of the receiver.*

The proof is in Appendix A.2.

To gain the intuition for Proposition 3 and the role of assumption (2), notice that an immediate consequence of (2) is that the expected utilities from any outcome μ satisfy $V_2(\mu) = b - cV_1(\mu)$. Thus, for any two outcomes μ' and μ''

$$V_1(\mu') \geq (>) V_1(\mu'') \iff V_2(\mu') \leq (<) V_1(\mu''). \tag{3}$$

In words, assumption (2) generalizes zero-sum preferences and implies that there is no room for cooperation: what is better for one is always worse for the other.

Now consider sequential obfuscation (the argument for sequential disclosure is analogous). By (3), for every outcome μ, at least one sender prefers no disclosure to μ. As no disclosure can be enforced individually by each sender, it must be an equilibrium outcome of the sequential obfuscation game.

Remark 1. In the zero-sum-like situation stipulated by assumption (2), one could expect that the second mover has an advantage. Curiously, as apparent from Proposition 3, this need not be the case in sequential obfuscation and sequential disclosure games. Regardless of the order of moves, the sender who prefers no disclosure always wins in sequential obfuscation, and the sender who prefers full disclosure always wins in sequential disclosure.

3.3. **Risk averse senders.** We now consider the case that is particularly relevant for applications. We assume that the senders are risk averse, that is and both u_1 and u_2 are weakly concave. We show that under this assumption sequential obfuscation leads to no disclosure. We thus obtain a clearcut comparison between sequential disclosure and sequential obfuscation in this case.

Proposition 4. *Suppose that both senders are risk averse. Then no disclosure is an equilibrium outcome of sequential obfuscation. Moreover, no disclosure is the unique equilibrium outcome if at least one sender is strictly risk averse, or for a generic decision function $a(\cdot)$ of the receiver.*

The proof is in Appendix A.3.

The intuition behind Proposition 4 is similar to that of Proposition 3. When both senders are risk averse, for any outcome μ, at least one of the senders prefers no disclosure to μ. In sequential obfuscation no disclosure can be enforced individually by each sender, so every other outcome will be "blocked" by one of the senders (unless it is as good as no disclosure for both of them, which generically does not occur).

Unlike sequential obfuscation, sequential disclosure does not lead to no disclosure when both senders are risk averse. In fact, by Corollary 1, sequential disclosure leads to revelation of some information (except when both senders prefer no disclosure to all other outcomes), and it can even lead to full disclosure. For example, let

$$u_1(y) = \sqrt{y},\ u_2(y) = \sqrt{1-y},\ \text{and}\ a(x) = 1 - e^{-x}.$$

Then $u_1(a(x)) = \sqrt{1 - e^{-x}}$ is strictly concave in x and $u_2(a(x)) = \sqrt{e^{-x}}$ is strictly convex in x. This means that the unique most preferred outcome of sender 1 is no disclosure and the unique most preferred outcome of sender 2 is full disclosure (e.g., Kamenica and Gentzkow, 2011). By Proposition 4, sequential obfuscation leads to no disclosure. In contrast, in the sequential disclosure game, full disclosure is the unique

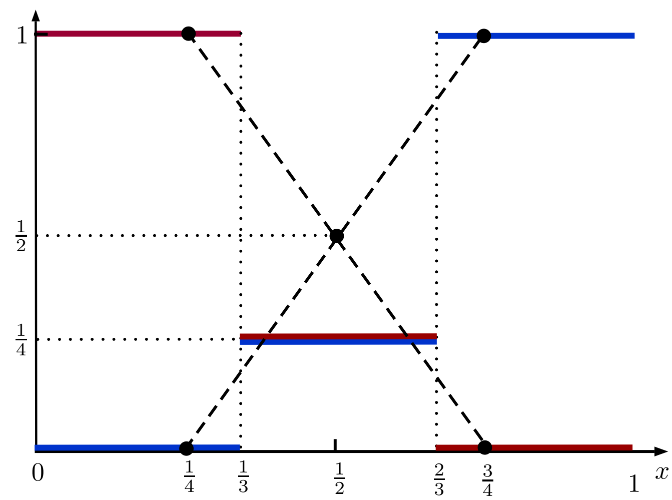

FIGURE 1. Utility $u_1(a(x))$ of sender 1 (solid blue) and utility $u_2(a(x))$ of sender 2 (solid red).

unimprovable outcome for sender 2. Thus, by Proposition 1, sequential disclosure leads to full disclosure.

Interestingly, a symmetric claim to Proposition 4, that if both senders are risk seeking then sequential disclosure fully reveals the state, need not be true. It is only true in the case of two states, that is, when the prior μ_0 has support $\{0, 1\}$. For a counterexample, let μ_0 be uniform on $[0, 1]$, and let

$$u_1(y) = y^2,\ u_2(y) = (1-y)^2, \text{ and } a(x) = \begin{cases} 0 & \text{if } x \in [0, 1/3], \\ 1/2 & \text{if } x \in (1/3, 2/3), \\ 1 & \text{if } x \in [2/3, 1]. \end{cases}$$

Then $u_1(a(x))$ and $u_2(a(x))$ are as shown in Fig. 1. Let us compare the full disclosure and the cutoff disclosure $\mu_{1/2}$ that reveals whether the state is above or below 1/2. Observe that $\mu_{1/2}$ induces the posteriors 1/4 and 3/4 equally likely, and yields the expected utility of 1/2 for both senders (illustrated by the midpoint of dashed lines in Fig. 1). However, full disclosure yields the expected utilities

$$\int_0^1 u_i(a(x))\mathrm{d}x = \frac{1}{3}\cdot 0 + \frac{1}{3}\cdot\frac{1}{4} + \frac{1}{3}\cdot 1 = \frac{5}{12} < \frac{1}{2}, \quad i = 1, 2.$$

That is, both senders strictly prefer $\mu_{1/2}$ to full disclosure. Consequently, by Proposition 1, full disclosure cannot be an equilibrium outcome of sequential disclosure, even though both senders are risk seeking.

Appendix

A.1. **Proof of Proposition 2.** Suppose by contradiction that $\mu^O \succeq \mu^D$. Then μ^O is attainable by disclosure from μ^D, and μ^D is attainable by obfuscation from μ^O. Because $\mu^D \in \mathcal{M}_2^D$, we have $V_2(\mu^D) \geq V_2(\mu^O)$. Because $\mu^O \in \mathcal{M}_2^O$, we have $V_2(\mu^O) \geq V_2(\mu^D)$. Hence

$$V_2(\mu^D) = V_2(\mu^O),$$

so sender 2 is indifferent between μ^D and μ^O.

Next, μ^O must be unimprovable by disclosure, so $\mu^O \in \mathcal{M}_2^D$. Indeed, if μ^O was improvable by disclosure to some μ, so $\mu \succeq \mu_O$, then μ^D would have been improvable by disclosure to μ as well, because $\mu \succeq \mu^O \succeq \mu^D$. Similarly, $\mu^D \in \mathcal{M}_2^O$. We thus obtain that both μ^D and μ^O are feasible choices for sender 1 in both problems.

Because μ^D is an equilibrium outcome in sequential disclosure but μ_O is feasible for sender 1, by Proposition 1 we have $V_1(\mu^D) \geq V_1(\mu^O)$. Analogously, because μ^O is an equilibrium outcome in sequential obfuscation but μ_D is feasible for sender 1, we have $V_1(\mu^O) \geq V_1(\mu^D)$. Hence

$$V_1(\mu^D) = V_1(\mu^O),$$

so sender 1 is indifferent between μ^D and μ^O. We thus have reached a contradiction to the assumption that $\big(V_1(\mu^D), V_2(\mu^D)\big) \neq \big(V_1(\mu^O), V_2(\mu^O)\big)$. $\square$

A.2. **Proof of Proposition 3.** Consider sequential obfuscation (the proof for sequential disclosure is analogous). Denote by μ^{ND} the no disclosure outcome. Each outcome μ in $\mathcal{M}_2^O$ is unimprovable by obfuscation for sender 2, in particular it is unimprovable by no disclosure. Thus we have

$$V_2(\mu) \geq V_2(\mu^{ND}) \text{ for each } \mu \in \mathcal{M}_2^O.$$

Then by (3) we have

$$V_1(\mu) \leq V_1(\mu^{ND}) \text{ for each } \mu \in \mathcal{M}_2^O.$$

It follows from Proposition 1 that no disclosure is an equilibrium outcome of sequential obfuscation.

To show the uniqueness for nongeneric $a(\cdot)$, observe that it follows from Proposition 1 that if two outcomes μ' and μ'' that satisfy $\mu' \succeq \mu''$ are both equilibria, then (3) must hold as equality for both senders. But because $\mu' \succeq \mu''$, (3) cannot hold as equality for a generic $a(\cdot)$ unless $\mu' = \mu''$. $\square$

A.3. **Proof of Proposition 4.** Let $\mu \in \mathcal{M}$ be an equilibrium outcome of sequential obfuscation. Let x_0 be the prior expected state, $x_0 = \int_{x\in\Theta} x \mathrm{d}\mu_0(x)$. Denote by μ^{ND} the no disclosure outcome. This outcome assigns probability 1 to x_0.

Let a^{-1} be the (generalized) inverse of a. Let $y = a(x)$. Then $\mu(a^{-1}(y))$ is the probability distribution of y over the domain $\tilde{A} = [a^{-1}(0), a^{-1}(1)]$. For each $i = 1, 2$, by the concavity of u_i and Jensen's inequality we have

$$\begin{aligned} u_i\left(\int_{x\in\Theta} a(x)\mathrm{d}\mu(x)\right) &= u_i\left(\int_{y\in\tilde{A}} y\mathrm{d}\mu(a^{-1}(y))\right) \\ &\geq \int_{y\in\tilde{A}} u_i(y)\mathrm{d}\mu(a^{-1}(y)) = \int_{x\in\Theta} u_i(a(x))\mathrm{d}\mu(x). \end{aligned} \tag{4}$$

Next, using the assumption that μ is an equilibrium outcome of sequential obfuscation, for each $i = 1, 2$ we have

$$\int_{x\in\Theta} u_i(a(x))\mathrm{d}\mu(x) \geq u_i(a(x_0)). \tag{5}$$

Inequality (5) holds for $i = 2$ because by Proposition 1 we have $\mu \in \mathcal{M}_2^O$, and thus μ is unimprovable by μ^{ND}. To establish inequality (5) for $i = 1$, note that $\mu^{ND} \in \mathcal{M}_2^O$, because no disclosure μ^{ND} is trivially unimprovable by obfuscation. Thus, by Proposition 1 sender 1 must weakly prefer μ to μ^{ND}.

By (1), u_1 is strictly increasing and u_2 is strictly decreasing. Thus (4) and (5) imply

$$u_i\left(\int_{x\in\Theta} a(x)\mathrm{d}\mu(x)\right) = u_i(a(x_0)) \quad \text{for each } i = 1, 2. \tag{6}$$

It follows that either $\mu = \mu^{ND}$, or the senders are indifferent between μ and μ^{ND}, so both μ and μ^{ND} are equilibrium outcomes.

Lastly, suppose that $\mu \neq \mu^{ND}$. Because μ is a mean-reserving spread of μ^{ND}, equation (6) cannot be satisfied for a generic $a(\cdot)$, or when either u_1 or u_2 are strictly concave. $\square$

References

Alonso, R., and O. Câmara (2016): "Political Disagreement and Information in Elections," *Games and Economic Behavior*, 100, 390–412.

Ambrus, A., and S. E. Lu (2014): "Almost fully revealing cheap talk with imperfectly informed senders," *Games and Economic Behavior*, 88, 174–189.

Ambrus, A., and S. Takahashi (2008): "Multi-sender cheap talk with restricted state spaces," *Theoretical Economics*, 3, 1–27.

Arieli, I., Y. Babichenko, R. Smorodinsky, and T. Yamashita (2022): "Optimal persuasion via bi-pooling," *Theoretical Economics*, forthcoming.

Au, P. H., and K. Kawai (2020): "Competitive information disclosure by multiple senders," *Games and Economic Behavior*, 119, 56–78.

Austen-Smith, D. (1993): "Interested experts and policy advice: Multiple referrals under open rule," *Games and Economic Behavior*, 5, 3–43.

Battaglini, M. (2002): "Multiple referrals and multidimensional cheap talk," *Econometrica*, 70, 1379–1401.

——— (2004): "Policy advice with imperfectly informed experts," *The BE Journal of Theoretical Economics*, 4, Article 1.

Blackwell, D. (1953): "Equivalent Comparisons of Experiments," *Annals of Mathematical Statistics*, 24, 265–272.

Bloedel, A. W., and I. Segal (2020): "Persuading a Rationally Inattentive Agent," mimeo.

Boleslavsky, R., and C. Cotton (2018): "Limited capacity in project selection: Competition through evidence production," *Economic Theory*, 65, 385–421.

Chan, J., S. Gupta, F. Li, and Y. Wang (2019): "Pivotal persuasion," *Journal of Economic theory*, 180, 178–202.

D'Agostino, E., and D. J. Seidmann (2022): "The order of presentation in trials: Plaintive plaintiffs," *Games and Economic Behavior*, 132, 328–336.

Degan, A., and M. Li (2016): "Persuasion with costly signals," mimeo.

Duffie, D., P. Dworczak, and H. Zhu (2017): "Benchmarks in Search Markets," *Journal of Finance*, 72, 1983–2044.

Dworczak, P., and G. Martini (2019): "The Simple Economics of Optimal Persuasion," *Journal of Political Economy*, 127, 993–2048.

Edmond, C., and Y. K. Lu (2021): "Creating confusion," *Journal of Economic Theory*, 191, 105145.

Gehlbach, S., and K. Sonin (2014): "Government Control of the Media," *Journal of Public Economics*, 118, 163–171.

Gentzkow, M., and E. Kamenica (2016): "A Rothschild-Stiglitz Approach to Bayesian Persuasion," *American Economic Review: Papers & Proceedings*, 106, 597–601.

——— (2017a): "Bayesian persuasion with multiple senders and rich signal spaces," *Games and Economic Behavior*, 104, 411–429.

——— (2017b): "Competition in Persuasion," *Review of Economic Studies*, 85, 300–322.

Gilligan, T. W., and K. Krehbiel (1989): "Asymmetric information and legislative rules with a heterogeneous committee," *American Journal of Political Science*, 33, 459–490.

Ginzburg, B. (2019): "Optimal Information Censorship," *Journal of Economic Behavior and Organization*, 163, 377–385.

Gitmez, A., and P. Molavi (2020): "Media Capture: A Bayesian Persuasion Approach," mimeo.

Goldstein, I., and Y. Leitner (2018): "Stress Tests and Information Disclosure," *Journal of Economic Theory*, 177, 34–69.

Hedlund, J. (2017): "Bayesian persuasion by a privately informed sender," *Journal of Economic Theory*, 167, 229–268.

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

Kleiner, A., B. Moldovanu, and P. Strack (2021): "Extreme points and majorization: Economic applications," *Econometrica*, 89, 1557–1593.

Koessler, F., and V. Skreta (2021): "Information design by an informed designer," CEPR Discussion Paper No. DP15709.

Kolotilin, A. (2015): "Experimental Design to Persuade," *Games and Economic Behavior*, 90, 215–226.

——— (2018): "Optimal information disclosure: A linear programming approach," *Theoretical Economics*, 13, 607–636.

Kolotilin, A., T. Mylovanov, and A. Zapechelnyuk (2022): "Censorship as optimal persuasion," *Theoretical Economics*, 17, 561–585.

Kolotilin, A., T. Mylovanov, A. Zapechelnyuk, and M. Li (2017): "Persuasion of a Privately Informed Receiver," *Econometrica*, 85, 1949–1964.

Kolotilin, A., and A. Zapechelnyuk (2019): "Persuasion Meets Delegation," mimeo.

Krishna, V., and J. Morgan (2001a): "Asymmetric information and legislative rules: Some amendments," *American Political Science Review*, 95, 435–452.

——— (2001b): "A model of expertise," *Quarterly Journal of Economics*, 116, 747–775.

LEVY, G., AND R. RAZIN (2007): "On the limits of communication in multidimensional cheap talk: A comment," *Econometrica*, 75, 885–893.

LI, F., AND P. NORMAN (2018): "On Bayesian persuasion with multiple senders," *Economics Letters*, 170, 66–70.

——— (2021): "Sequential persuasion," *Theoretical Economics*, 16, 639–675.

LI, F., Y. SONG, AND M. ZHAO (2022): "Global manipulation by local obfuscation," Available at SSRN 3471491.

LI, M. (2010): "Advice from multiple experts: A comparison of simultaneous, sequential, and hierarchical communication," *The BE Journal of Theoretical Economics*, 10, Article 18.

LIPNOWSKI, E., L. MATHEVET, AND D. WEI (2020): "Attention Management," *American Economic Review: Insights*, 2, 17–32.

MYLOVANOV, T., AND A. ZAPECHELNYUK (2013a): "Decision rules revealing commonly known events," *Economics Letters*, 119, 8–10.

——— (2013b): "Optimal arbitration," *International Economic Review*, 54, 769–785.

——— (2022): "A model of debates: Moderation vs free speech," mimeo.

ORLOV, D., P. ZRYUMOV, AND A. SKRZYPACH (2020): "Design of Macro-prudential Stress Tests," mimeo.

PEREZ-RICHET, E. (2014): "Interim bayesian persuasion: First steps," *American Economic Review*, 104, 469–74.

RAVINDRAN, D., AND Z. CUI (2020): "Competing persuaders in zero-sum games," working paper, arXiv 2008.08517.

ROMANYUK, G., AND A. SMOLIN (2019): "Cream Skimming and Information Design in Matching Markets," *American Economic Journal: Microeconomics*, 11, 250–276.

WOLINSKY, A. (2002): "Eliciting information from multiple experts," *Games and Economic Behavior*, 41, 141–160.

ZAPECHELNYUK, A. (2013): "Eliciting information from a committee," *Journal of Economic Theory*, 148, 2049–2067.

——— (2020): "Optimal Quality Certification," *American Economic Review: Insights*, 2, 161–176.

——— (2022): "On the equivalence of information design by uninformed and informed principals," SSRN: http://dx.doi.org/10.2139/ssrn.4017510.