Persuasion under “Aspect-Restricted” Investigation and Strategic Communication*

Antonio Jiménez-Martínez†

December, 2019

Abstract

This article explores how a Sender can persuade a Receiver to accept a proposal under two key considerations. The relevant uncertainty has two distinct *aspects*. The first consideration is that the Sender can design information—or, equivalently, make (ex ante) *investigation* decisions—over any single one of the two aspects. Secondly, while the Sender knows the true joint prior over the two aspects, the Receiver knows the marginal prior over each of the two *separate* aspects but is uncertain about the dependencies between the aspects. The Sender can then complement the information disclosed by the selected information structure with (interim) communication about the relationships between the two aspects. In equilibrium, the Sender selects one of the aspects, and optimally designs information over it, depending on the marginal priors over the separate aspects and on the Receiver's preferences. A salient class of equilibria features full revelation of the Sender's private information about the relationships between the two aspects.

JEL Classification: D72; D82; D83

*I thank Nageeb Ali, Emir Kamenica, Karna Basu, Héctor Chade, David Dillenberger, Mauricio Fernández, Simone Galperti, Matthew Gentzkow, Johannes Hörner, Navin Kartik, Andreas Kleiner, Isabel Melguizo, Stephen Morris, Charles Sprenger, and Philipp Strack for very insightful conversations and helpful comments. I gratefully acknowledge financial support from CONACYT through SNI Fellowship 41826.

†CIDE, Division of Economics, Mexico City, MEXICO. Email: antonio.jimenez@cide.edu.

1. Introduction

Investigation decisions are central to influence decision making.[1] The literature on information design[2] (Bergemann and Morris, 2013, 2016, 2019; Taneva, 2018) formalizes investigation decisions as the (ex ante) choice of information structures that determine the information that will be subsequently disclosed. Information design has explored influential communication under the key assumption that the Sender has *full (and complete) commitment* power over all the dimensions, or *aspects*, of the relevant uncertainty:[3] the Sender can commit to provide the Receiver with any rule that maps the *entire* state of the world into a probability distribution over action recommendations.[4]

For environments where the commitment assumption is less natural, the canonical cheap talk (Green and Stokey, 1980; Crawford and Sobel, 1982) and verifiable disclosure (Grossman, 1981; Milgrom, 1981; Milgrom and Roberts, 1986) models of (interim) communication consider instead that, after learning the true realization of the state—their *types*—, Senders decide how much to communicate by anticipating the Receivers' induced actions.[5]

In practice, many multi-dimensional-uncertainty situations feature combinations of (ex ante) information design and (interim) communication. This paper's approach rests on the consideration that, due perhaps to high numbers of relevant aspects of uncertainty, to strong conceptual distinctions between different aspects, to technological bounds on investigation, or to information-processing restrictions,[6] investigation capable of disclosing information over

[1] For concreteness, I use the term *investigation* to refer broadly to scientific research, journalistic investigation, fact-based studies, external expert consultation, audits, polls, trials, medical tests, or experimentation.

[2] This approach is also known also as Bayesian persuasion (Rayo and Segal, 2010; Kamenica and Gentzkow, 2011; Gentzkow and Kamenica, 2017a,b) for situations with a single Sender and a single Receiver with no private information. The current paper builds more closely upon the developments on information design and thus relies more on its terminology.

[3] Following the pertinent literature, the term *full* refers to that commitment is entirely binding for the Sender. For those cases where there is some probability that commitment is not binding, the information design literature uses the terms *partial* or *limited* commitment. For multi-dimensional environments, I reserve the term *complete* to consider that commitment power is full over all the dimensions, whereas I propose the term "aspect-restricted" commitment to describe situations where commitment power is also full but only available over a (strict) subset of the dimensions.

[4] Using a Revelation Principle argument, this is equivalent to map the state of the world into a probability distribution over signals which, in turn, provide new information for the Receivers to choose their preferred actions. Using signals instead of action recommendations is the formulation most typical in Bayesian persuasion models.

[5] Relative to information design, the lack of commitment under cheap talk makes influential communication more limited.

[6] Consider, e.g., restrictions to sample size, high costs of conducting research, time constraints for decision

all aspects is not always feasible on time before the decision making is due. Conceivably, sometimes Senders have commitment power only over a subset of all the relevant aspects before the decision must be made. In these situations, Senders often tend to complement the information disclosed by the information structures—selected over a subset of all the relevant aspects—with (interim) communication, or "interpretative views,"—sometimes using unverifiable messages, other times using verifiable evidence—about how such aspects relate to the remaining aspects—over which there was no information design. Arguably, cheap talk—or verifiable disclosure—communication requires less effort and time than making choices about investigation, and then obtaining and analyzing the results from the selected investigation processes. In a stylized fashion, this paper aims at capturing these situations by considering that the relevant uncertainty has two aspects and that the Sender can commit for information design over any single one of them before the decision must be made. In addition, the Sender, in his position as an expert, is assumed to have an informational advantage over the Receiver about the dependencies between the two aspects. Given this, the Sender can also make information revelation decisions about his private knowledge of the dependencies between the two aspects. To capture the latter feature, I make for simplicity the extreme assumption that the Sender is fully informed about the dependencies between the two aspects, whereas the Receiver is uncertain about such dependencies.

The proposed benchmark aims at providing foundations to study how investigation choices over a subset of the relevant aspects interact with "interpretative" communication about the influence of the investigated aspects on the remaining ones. The analysis focuses on situations where a Sender wants to persuade a Receiver that must decide between accepting or rejecting a proposal. The approach is abstract but the underlying ideas are of substantial relevance. To fix ideas, consider the consulting CEO (Sender/expert) of an automobile company who advises the stock-holders Board of the firm (Receiver/decision-maker). The Board must decide whether or not to launch the company's flagship electric model into a new market. The relevant state of the world $\theta = (x, y)$ consists of two aspects about the potential customers of this new market: their future income (x) and their environmental concerns (y). The CEO wishes the firm to launch the new model regardless of the true state of the world, whereas the Board wishes to do so only if the state meets certain conditions. Although investigation can usually be conducted separately about income growth and about environmental concerns, it seems less feasible that a common investigation process be able to disclose information about both aspects combined. This could be caused by the two aspects

making, or limits on the amount of information that can be accumulated or processed by Receivers.

being conceptually very different. The proposed model assumes first that the Board can use the data from investigation about any single one of the two aspects before her decision is due. I reserve the term *aspect-restricted commitment* to describe situations where the expert cannot select information structures about the pair (x, y), neither about the correlations between x and y. This model assumes that the expert is restricted to selecting information structures for any of the two variables, income growth (x) or environmental concerns (y), *separately*. In addition, the expert can select only one of them before the decision-maker is due to decide whether or not to launch the new model. The second key assumption of the model is that the expert has an informational advantage over the decision-maker about the actual relationships between the two aspects. Suppose that the expert and the decision-maker (commonly) know the marginal priors over the variables x and y, but that the dependencies between x and y are not known by the decision-maker, according to her priors. Formally, the decision-maker is assumed to have *non-fully identified* priors, where the missing piece of information is how each aspect depends on the other. On the other hand, on the basis of his position as informed expert, the Sender is assumed to have his own private information—perhaps unverifiable or in the form of previously collected previously verifiable reports—about how income growth relates to environmental concerns. To consider a setup as much stylized as possible, I assume that the expert has in fact full information about such dependencies.[7] Then, alongside with his investigation choice, the expert can also choose how much information he reveals about the relationship between income growth and environmental concerns.

The model accommodates both ways of costless (interim) communication about the relationships between the aspects, either through cheap talk or through verifiable messages. Investigation over one of the aspects together with information about how the two aspects depend on each other does indeed provide new information about the two aspects of uncertainty. Since I am interested in exploring how ex ante information provision interacts with interim communication, assuming that the Sender is able to choose investigation over only one of the two aspects is a convenient analytical condition to address this research question.

While the focus on a two-action setting seems natural for many environments, it also stands as a key analytical requirement if we wish to explore how information design over the separate aspects interacts with strategic communication about their relationships. If

[7] One could alternatively consider that the expert has instead noisy signals about such relationships so that he is "better informed" than the decision-maker, yet not fully informed. Nonetheless, this would complicate unnecessarily the analysis as the model's qualitative results would continue to hold.

we considered more than two actions in the proposed setup, then the Receiver could use the Sender's preferred investigation choices to learn how the distinct aspects relate to each other. Since we wish to avoid such a contaminating effect in information disclosure, the model assumes a two-action setting where one of the actions is always preferred by the Sender—regardless of the state of the world. The conflict of interests between the two actors is assumed to be high. Based on any possible priors about the state of the world, the Sender always prefers acceptance, whereas the Receiver wishes rejection.

A key contribution of the paper is to identify the features of priors about the state and of the players' preferences that characterize the Sender's optimal disclosure, taking into account that, in equilibrium, the Receiver both obeys the recommendations from the (ex ante) investigation choices and best responds to the (interim) communication strategies.

Suppose first that the Sender chooses to fully reveal his private information about the dependencies between the two aspects. Then, under the assumed restrictions on how the Sender can make investigation choices, he first considers optimal investigation choices (separately) for each of the two aspects. It turns out that investigation over one aspect is able to persuade the Receiver only if, in spite of the sharp conflict of interests assumed, there is at least one realization of such an aspect over which the players' interests are indeed perfectly aligned. If such a type of aspect realizations exist, we say that they constitute an *agreement set* over the corresponding aspect. Going back to our example, suppose that for each of the two aspects, income growth and environmental concern, investigation can release at least one piece of data (over the corresponding aspect) upon which the Board wishes to launch the new model. If that is the case, the CEO will then consider investigation processes that recommend acceptance with probability one for possible growth rates, on one side, and for possible environmental concerns indexes, on the other side, that induce the Board to accept the proposal. In addition, provided that the agreement sets over the aspects are nonempty, the optimally selected investigations will recommend also acceptance with some probability (between zero and one) also for other realizations, outside of the corresponding agreement set. Such recommendations for realizations outside of the agreement set (for the corresponding aspect) will be optimally made according to a ranking that captures the degree to which the Board is relatively more willing to launch the new model—i.e., the degree to which the players' interests are (partially) aligned. Given this, using backwards induction reasoning, the CEO compares then what investigation is able to attain for income growth and for environmental concerns and selects the aspect that gives him the higher expected utility. In short, the Sender will select the aspect such that the possible realizations of the aspect in average alleviates the most the original source of conflict of interests. The optimal aspect choice, as well as the chosen investigation processes, are "tailored" to the marginal priors about each aspect and, crucially, to the Receiver's tastes as well.

As to whether the Sender in fact wants to reveal his private information about how the two aspects correlate, it turns out that aspect-restricted information design interacts in an interesting way with strategic communication. I make the common assumption that the Receiver is maximally skeptical when the Sender's communication strategy does not allow to learn the Sender's type. Furthermore, when maximal skepticism is not enough to single out one possible type of the Sender as more likely than the others, I assume that the Receiver updates her beliefs about how the aspects correlate in a way proportional to her initial priors. This seems the most neutral consideration that respects Bayesian updating under maximal skepticism. Under these assumptions, any equilibrium disclosure strategy features full revelation of the Sender's private information, regardless of whether communication is made through cheap talk or through verifiable messages.

The logic behind the full revelation result lies in that, under the proposed two-action setting, information design over a single aspect disciplines the Receiver in a way such that she is left indifferent between any subset of different types that are pooled together. Then, in spite of her skepticism, the Receiver places positive probability on each type from any subset of pooled types. In addition to this, the linearity of the information design problem crucially leads to that the optimal expected utility that the Sender receives, for a subset of pooled types, can be expressed as a convex combination of the optimal expected utilities that he would receive if he fully separated each of the types from the subset. As a consequence, the type that obtains the highest utility within the pooled subset has incentives to separate from the rest of the pooled types. Therefore, on the one hand, (ex ante) investigation alleviates the initial conflict of interests between the parties over the possible relationships between the two aspects. On the other hand, investigation allows the Sender to receive higher utility when he is able to credibly communicate some possible priors about the state. Intuitively, for such priors about the state, investigation is able to put relatively high probabilities on the Receiver wishing to accept to proposal. The full revelation mechanism that this paper obtains crucially relies on the combination of the incentive-compatibility constraint for the (ex ante) information design problem with the way in which the skeptical Receiver updates her priors and the Sender accordingly best replies. This mechanism is, therefore, rather different from the classical "unravelling" argument of the verifiable disclosure models as information design makes the Receiver indifferent between any subset of different types pooled by the Sender. In particular, the Receiver always has a ranking of types that conflicts with that of the Sender in verifiable disclosure models.

The next Section 2 outlines the model. The leading example illustrates, in Subsection 2.6, the model and the logic of the proposed equilibrium. Section 3 introduces formally the equilibrium concept and compares it to the traditional information design framework under full commitment. The main results are presented in Section 4. Section 5 applies the model to provide a plausible interpretation and rationale for slant in media persuasion. The closest related literature is discussed in Section 6, and Section 7 concludes. All the proofs are relegated to the Appendix.

2. Model

I will use $\tilde{z}$ to denote a random variable with realization z and $\mathbb{E}_\xi[\cdot]$ to indicate the expected value with respect to probability distribution ξ.

There are two players, indexed by $i = S, R$, a partially informed *(S)ender* (he) and an uninformed *(R)eceiver* (she). The model considers two different channels for disclosing information. "Investigation over a given aspect of uncertainty" will be modeled following the information design/Bayesian persuasion approach—e.g., Kamenica and Gentzkow (2011) and the literature reviewed by Bergemann and Morris (2019). "Communication about the relationships between the aspects" will be modeled as in the classical cheap talk—e.g., Green and Stokey (1980) and Crawford and Sobel (1982) —and persuasion/verifiable disclosure —e.g., Grossman (1981), Milgrom (1981), and Milgrom and Roberts (1986)—approaches. The difference between both ways of information transmission will then be based on what the Sender knowns when he decides how to disclose, and on his ability to commit to his disclosure decision. The Sender commits ex ante in his investigation, or information design, decision, whereas he communicates at an interim stage (thus, conditional on his private information) in his strategic communication decision.

The Receiver takes an *action* a from a set $A \equiv \{\underline{a}, \overline{a}\}$. The low action $\underline{a}$ means *rejecting* a certain proposal and the high action $\overline{a}$ means *accepting* the proposal. The Receiver cares about action a and about a two-dimensional[8] *state of the world* $\theta \equiv (x, y) \in \Theta \equiv X \times Y \subset \mathbb{R}^2$. Each set $\mathcal{K} \in \{X, Y\}$ is finite and describes the possible realizations of the respective *aspect*

[8] The model is developed in terms of a two-dimensional state for simplicity. Its functioning and main implications, though, hold qualitatively for a general multi-dimensional state with a finite number of dimensions.

$\kappa \in \{x, y\}$ of the underlying state θ.[9] For a given aspect $\kappa \in \mathcal{K}$, let $-\kappa$ identify the remaining aspect—i.e., $\{-\kappa\} \equiv \{x, y\} \setminus \{\kappa\}$—and $-\mathcal{K}$ the remaining set of aspects—i.e., $\{-\mathcal{K}\} \equiv \{X, Y\} \setminus \{\mathcal{K}\}$. To simplify the exposition of the model, I assume that X and Y have the same cardinality $m \geq 3$,[10] with $X = \{x_1, \ldots, x_j, \ldots, x_m\}$ and $Y = \{y_1, \ldots, y_k, \ldots y_m\}$. When indexes are explicitly used, $\kappa_l \in \mathcal{K}$ will indicate the generic realization of an aspect, either $x_j \in X$ or $y_k \in Y$.

2.1. Initial Information Structure

The players begin with common priors over each (separate) aspect of uncertainty, which are given by (marginal) distributions $\psi_x \in \Delta_{++}(X)$ and $\psi_y \in \Delta_{++}(Y)$. In some parts of the paper, it will be useful to capture each marginal distribution ψ_κ by means of an m-dimensional vector $\psi_\kappa = (\psi_\kappa(\kappa_l))_{l=1}^m$. The relationships between aspects x and y are described by a set of (conditional) probability distributions that I label as a *pattern of dependence* τ. The logic behind the notion of pattern of dependence is as follows. First, let me use $\psi_\tau \in \Delta_{++}(\Theta)$ to denote the particular (joint) *prior about the state of the world* θ that corresponds to the pattern of dependence τ. Then, I will use $t_{xy} \equiv \psi_\tau(y \mid x)$ to denote the conditional probability, according to the prior ψ_τ, of realization y given x. In addition, when we consider realizations as $x = x_j$ and $y = y_k$, I shall sometimes use the short-hand notation $t_{jk} \equiv \psi_\tau(y_k \mid x_j)$ to work with matrix notation. Each m-dimensional vector $t_j \equiv (t_{jk})_{k=1}^m$ (for $j = 1, \ldots, m$) of probabilities characterizes a (conditional) probability distribution $t_j \in \Delta(Y)$ over aspect y. Using matrix notation, the notion of *pattern of dependence* between the two aspects of uncertainty can then be algebraically captured by an $m \times m$ matrix[11]

$$\tau \equiv [t_1 \ \cdots \ t_j \ \cdots \ t_m].$$

Conceptually, a pattern of dependence τ describes how each realization $y_k \in Y$ depends on each realization $x_j \in X$.[12] Let $\mathcal{T}_j \subset \Delta(Y)$ be the set of all possible conditional distributions

[9] With a slight abuse of notation, in some parts of the paper, κ will denote both dimensions and realizations.

[10] We need to require more than two possible realizations of each aspect for technical reasons that will be explained in Subsection 2.1.

[11] As usual, the vectors t_j are considered as column vectors.

[12] Note that I am choosing without loss of generality aspect x as reference to describe the dependencies between the two aspects. In particular, Bayesian consistency requires that each pair of conditional probabilities $\psi_\tau(y \mid x)$ and $\psi_\tau(x \mid y)$ be related through the condition: $\psi_\tau(y \mid x)\psi_x(x) = \psi_\tau(x \mid y)\psi_y(y)$ for each $\theta = (x, y) \in \Theta$.

t_j and let $\mathcal{T} \equiv \times_{j=1}^{m} \mathcal{T}_j$ be the set of all possible patterns of dependence τ. The model considers a finite set $\mathcal{T} = \{\tau_1, \tau_2, \ldots, \tau_s\}$ of possible patterns of dependence, with $s \geq 2$.

A key plausibility requirement arises in the proposed environment: marginal distributions/beliefs about the separate aspects and patterns of dependence must be consistent with each other. In particular, given a pattern of dependence $\tau \in \mathcal{T}$, the Radon-Nikodym Theorem allows us to use the definition of conditional probability to obtain that $\psi_\tau(\theta) = \psi_\tau(y \mid x)\psi_x(x) = t_{xy}\psi_x(x)$ for each $\theta = (x, y) \in \Theta$.[13] As a consequence, the notion of conditional probability imposes the following Bayesian plausibility condition on the marginal distributions of each of the two aspects and on the patterns of dependence $\tau \in \mathcal{T}$:

$$\sum_{x \in X} t_{xy}\psi_x(x) = \psi_y(y) \qquad \forall y \in Y.$$

This crucial Bayesian plausibility condition can be expressed more compactly using the matrix notation that I introduced above.

Assumption 1. *For any given pair of marginal distributions $\psi_x \in \Delta_{++}(X)$ and $\psi_y \in \Delta_{++}(Y)$ over the aspects of uncertainty, each possible pattern of dependence $\tau \in \mathcal{T}$ must satisfy the Bayesian plausibility condition: $\tau \cdot \psi_x = \psi_y$.*

The analysis considers that the available patterns of dependence τ that satisfy Assumption 1 above are not affected by the information design choices of the Sender. In other words, there is no feasible investigation about how the two relevant aspects of uncertainty relate to each other.

At this point I can comment on why $m \geq 3$ is a technically necessary condition in the proposed framework. First, it is needed to prevent the Receiver from learning the joint distribution ψ_τ by using the marginal distributions ψ_x and ψ_y. Secondly, if $m = 2$, then the system of linear equations $\tau \cdot \psi_x = \psi_y$ required by Assumption 1 is satisfied by a unique matrix τ. In that case, we would not be able to work with the key feature of the model that different patterns of dependence τ be Bayes−consistent with the marginal priors over the two aspects of uncertainty. When $m \geq 3$, the system $\tau \cdot \psi_x = \psi_y$ is undetermined and, therefore, it allows for multiple patterns of dependence τ to satisfy the required Bayesian plausibility condition, which suits nicely the proposed approach.

[13] Note that such a requirement is consistent with well-defined joint distributions over the overall state of the world. Specifically, $\sum_{\theta \in \Theta} \psi_\tau(\theta) = \sum_{x \in X} \psi_x(x) \sum_{y \in Y} t_{xy} = 1$ for each pattern of dependence $\tau \in \mathcal{T}$.

2.2. Preferences

The preferences of each player $i = S, R$ are described by an (ex post) utility function $u_i : A \times \Theta \to \mathbb{R}$. Each player i has a unique *ideal action* $a_i^*(\theta) \in A$ for each state realization, $\theta \in \Theta$, as well as a unique ideal action $\hat{a}_i(\tau) \in A$ for each possible prior over states $\psi_\tau \in \Delta_{++}(\Theta)$, for $\tau \in \mathcal{T}$. The players disagree on their ideal actions. The Sender strictly prefers acceptance regardless of the state of the world, whereas the Receiver only prefers acceptance if the state of the world belongs to a certain (nonempty) *acceptance set* $\overline{\Theta} \subset \Theta$. The model thus assumes a form of conflict of interests which is typically present in information design (under complete commitment), cheap talk, and verifiable disclosure environments.

Assumption 2. *The preferences of the players satisfy:*

(i) *each player $i = S, R$ has a unique ideal action $a_i^*(\theta) \equiv \arg\max_{a \in A} \; u_i(a, \theta)$ for each $\theta \in \Theta$;*

(ii) *(a) $u_S(\underline{a}, \theta) = u_R(\underline{a}, \theta) = 0$ for each $\theta \in \Theta$, and (b) $u_S(\overline{a}, \theta) > 0$ for each $\theta \in \Theta$, whereas $u_R(\overline{a}, \theta) > 0$ if $\theta \in \overline{\Theta}$ and $u_R(\overline{a}, \theta) < 0$ if $\theta \in \Theta \setminus \overline{\Theta}$;*

(iii) *each player $i = S, R$ has a unique ideal action $\hat{a}_i(\tau) \equiv \arg\max_{a \in A} \mathbb{E}_{\psi_\tau}[u_i(a, \tilde{\theta})]$ for each $\tau \in \mathcal{T}$;*

(iv) *$\hat{a}_R(\tau) = \underline{a}$ and $\hat{a}_S(\tau) = \overline{a}$ for each $\tau \in \mathcal{T}$;*

(v) *the Sender has a (strictly) monotone order over the set of patterns of dependence $\mathcal{T}$ according to his expected utility, provided that the Receiver accepts the proposal, that is, without loss of generality: $\mathbb{E}_{\psi_{\tau_1}}[u_S(\overline{a}, \tilde{\theta})] < \mathbb{E}_{\psi_{\tau_2}}[u_S(\overline{a}, \tilde{\theta})] < \cdots < \mathbb{E}_{\psi_{\tau_s}}[u_S(\overline{a}, \tilde{\theta})]$.*

Assumption 2 (iv) captures situations where the Sender wants to persuade the Receiver to move away from rejecting the proposal, while the latter would always reject in the absence of new information beyond the plausible priors about the underlying state. Assumption 2 (v) imposes a monotonicity condition on how the Sender ranks patterns of dependence, upon acceptance. This requirement has the flavor of the type of monotonicity typically considered in the classical cheap talk and persuasion models.

A central message of the literature on verifiable disclosure (Grossman, 1981; Milgrom, 1981; Milgrom and Roberts, 1986) is that when the preferences of the Sender are sufficiently opposed to those of the Receiver, then full disclosure is the unique equilibrium outcome. This insight stems from a classical "unravelling" argument supported by maximal skepticism on the Receiver's belief updating process when the Sender discloses truthful but incomplete information. Full disclosure also follows in the environment here explored. The mechanism underlying this result in the current paper, though, is quite different. Information design over one of the aspects disciplines the skeptical Receiver so as to leave her indifferent between any subset of types that are pooled by the Sender. Given this, the combination of aspect-restricted information design with (interim) communication is precisely what drives the full revelation result in the proposed setup.

2.3. Discussion of the Model

2.3.1. Novel Assumptions

The first novel element is that the Sender is constrained to choosing (fully committed) information structures over any single one of the two aspects. Investigation about the correlations between the two aspects is not feasible. Therefore, the design of information over the separate dimensions is assumed to leave unaffected the true pattern of correlation that relate the two aspects. One way to motivate these considerations—that I gather together under the term *aspect-restricted commitment*—rests on the observation that the separate dimensions of the uncertainty relevant for a decision problem often describe very conceptually distant features. Owing to such differences of substance, joint investigation over all dimensions is not always available, or feasible, in time before the Receiver is due to make her choice. Another way of motivate this assumption is by noting that, in many cases, the Receiver is not able to process and use the outcomes of a high number of separate investigation processes before the due date for her decision. Notwithstanding, information about the dependencies between the two aspects combined with investigation over a single aspect is able to disclose information over the two aspects.[14]

The second novel element is that the Sender privately knows the true relationships between the two aspects,[15] whereas the Receiver is uncertain about such dependencies. Then, starting from such a position of informational advantage, the Sender can resort to strategic (interim) communication about the dependencies between the aspects and, in this way, complement the information provided by the selected investigation.

[14] Using the semantics of information design, each (aspect-restricted) information structure over a single aspect, together with knowledge about how one aspect is related to another, induces a joint information structure over the two-dimensional state of the world.

[15] As commented in fn. 7, this assumption could be relaxed without substantial changes in the model's implications.

- Nature chooses a pattern of dependence and a state of the world
- Sender privately learns pattern of dependence
- Sender (i) selects one of the aspects and chooses investigation process over such an aspect, and (ii) communicates about the dependencies between the aspects
- Investigation discloses new information about the selected aspect
- Receiver updates beliefs based on the selected investigation and on received (interim) messages
- Receiver chooses action and both players obtain their payoffs

TABLE 1 Timing of the Information Design and Disclosure Game

2.3.2. Time Line

The timing of the disclosure game is as follows. First, Nature chooses: (i) a true pattern of dependence $\tau \in \mathcal{T}$ between the two aspects of uncertainty, and (ii) a true value of the state $\theta \in \Theta$ according to the selected prior $\psi_\tau \in \Delta_{++}(\Theta)$. The Sender learns privately the true pattern of dependence τ—his *type*. Secondly, without knowing the true value of the state θ, the Sender chooses a single aspect $\kappa \in \{x, y\}$ and an information structure over such a selected aspect—i.e., either over random variable $\tilde{x}$ or $\tilde{y}$. In addition, the Sender is allowed to complement the information disclosed by such an investigation choice with interim messages, either through cheap talk or through verifiable messages, about the dependencies between the two dimensions of the state—which are captured by his type τ . Unlike the decision on information design, the incentives of the Sender to communicate for each realization of τ are determined "in equilibrium"—exactly as in the classical cheap talk and disclosure/persuasion settings.

2.4. Communication about Pattern of Dependence

Nature selects a true pattern of dependence τ, which satisfies the Bayesian plausibility condition in Assumption 1, according to a (commonly known) *prior* $q \in Q \subseteq \Delta_{++}(\mathcal{T})$ *about patterns of dependence.*[16] The Sender privately learns the true realization of τ—his

[16] Accordingly, for each possible realization $x_j \in X$, each conditional distribution t_j is drawn from the set $\mathcal{T}_j$ according to a prior distribution $q_j \in Q_j \equiv \Delta_{++}(\mathcal{T}_j)$ where $q_j(t_j) = q(\tau)$ whenever the conditional distribution t_j is part of the pattern of dependence τ—i.e., whenever the vector t_j is included in matrix τ.

type—and, therefore, learns the true prior ψ_τ about the state θ. The Receiver, on the other hand, is initially uncertain about the possible patterns of dependence.[17]

The Sender decides how much information to convey to the Receiver about his type by selecting a *message* $d(\tau) \subseteq \mathcal{T}$ for each possible type τ. I focus on pure message strategies. A *message strategy* is a function $d : \mathcal{T} \rightarrow 2^{\mathcal{T}}$. This formulation is able to encompass both cheap talk communication (Green and Stokey, 1980; Crawford and Sobel, 1982) and verifiable disclosure models (Grossman, 1981; Milgrom, 1981; Milgrom and Roberts, 1986). To consider verifiable disclosure, we need to impose the condition that $\tau \in d(\tau)$ for each $\tau \in \mathcal{T}$, whereas, in principle, no further requirement is needed to capture cheap talk situations. Nonetheless, for cheap talk communication, recall that there are multiple message strategies d that lead to a common communication process about the Sender's type or, in other words, to a common system of induced posterior beliefs. I abstract from this typical multiplicity—which is due to the possible interpretations of messages—and focus on the class of message strategies where messages stick to their "literal meanings."[18] The model will restrict attention, both for cheap talk and for verifiable disclosure communication, to message strategies d such that (i) $\tau \in d(\tau)$ for each $\tau \in \mathcal{T}$, (ii) the Sender reveals fully his type τ by choosing $d(\tau) = \{\tau\}$, and (iii) the Sender pools, or withholds some information, over his type τ by choosing $d(\tau) = \mathcal{T}_\tau \neq \{\tau\}$ (for some non-singleton subset $\mathcal{T}_\tau \subset \mathcal{T}$). Let $\mathcal{D}$ be the set of all possible (pure) message strategies with "literal meanings" that satisfy these requirements. Furthermore, in some parts of the paper, we will be interested in paying special attention to message strategies where the Sender either pools over all his possible types—e.g., $d(\tau) = \mathcal{T}$ for each $\tau \in \mathcal{T}$—or reveals fully his true type—$d(\tau) = \{\tau\}$ for each $\tau \in \mathcal{T}$. For simplicity, I will then use the short-hand notation $d = w$ and $d = r$ to indicate that the Sender chooses, respectively, a (w)ithholding and a fully (r)evealing message strategy.

Suppose that the Sender selects a message strategy $d \in \mathcal{D}$. Then, conditional on the Sender's actual type being τ, the Receiver forms a posterior belief $\beta^d_\tau \in \Delta(\mathcal{T})$ about the pattern of dependence. Let $\beta^d \equiv \{\beta^d_\tau\}_{\tau \in \mathcal{T}}$ be a *system of posteriors induced by the message*

[17] In other words, for any possible type $\tau \in \mathcal{T}$ chosen by Nature, the Receiver begins with some *non-fully identified prior* $\zeta \in \Delta_{++}(\Delta_{++}(\Theta))$ about the state θ. While ζ can be thought of as a compound lottery over the set of states, ψ_τ gives us the simple lottery that corresponds to the pattern of dependence τ. In short, the Receiver has full information about the marginal distributions ψ_x and ψ_y of both aspects of the state but she is uncertain about the joint distribution ψ of the state, the missing piece of information being the pattern of dependence τ.

[18] This is without loss of generality since any possible pattern of communication can be specified using such a class of message strategies.

strategy d.[19] Furthermore, following also the approach of the classical persuasion literature, the model assumes that the Receiver's updating rule is based on maximal skepticism (or precaution). In particular, if the Sender pools over a subset of his possible types, then the Receiver is assumed to place probability one (or, in some cases, positive probability) on the particular type(s) that lead(s) her to the lowest possible expected payoffs, conditional on the optimal investigation choice followed by the Sender for each particular type. The specifics of the Receiver's skeptical Bayesian updating rule are detailed in Section 3, under Assumption 3.

2.5. Aspect-Restricted Information Design

Following the information design approach and its Revelation Principle arguments (Bergemann and Morris, 2019), the disclosure of information from investigation takes the form of direct "action recommendations."[20] The action recommendations provided by the selected information structures become public and cannot be subsequently concealed or distorted. Thus, the setup considers full commitment over each separate aspect of uncertainty.

Let me first review briefly a few key elements of the traditional information design approach to appreciate better how the proposed model builds upon the traditional framework—as well as how it differs from such a setup. The information design approach rests on the key concept of decision rule (Bergemann and Morris, 2013, 2016, 2019). For a given pattern of dependence $\tau \in \mathcal{T}$, a *decision rule* (under complete commitment) is a mapping $\rho^{\tau} : \Theta \to \Delta(A)$, where $\rho^{\tau}(a \mid \theta)$ is the probability according to which the Sender recommends the Receiver to choose action a if the true realization of the state is θ. Using the setup proposed in this paper so far, a decision rule (under complete commitment) ρ^{τ} would then satisfy the incentive-compatibility, or *obedience*, condition for a given pattern of dependence $\tau \in \mathcal{T}$ (i.e., under perfect information about τ) if

$$\sum_{\theta \in \Theta} \rho^{\tau}(\overline{a} \mid \theta)\ \psi_{\tau}(\theta)\ \left[u_R(\overline{a}, \theta) - u_R(\underline{a}, \theta)\right] \geq 0. \tag{1}$$

For a given prior ψ_{τ} (again, under perfect information about τ), the condition in Eq. (1) coincides formally with the obedience criterion required by Bergemann and Morris (2013,

[19] For example, for the fully revealing strategy $d = r$ it follows that $\beta^{r}_{\tau}(\tau) = 1$ and $\beta^{r}_{\tau}(\tau') = 0$ for each $\tau' \neq \tau$. On the other hand, for the withholding strategy $d = w$, it follows that the posterior beliefs β^{w}_{τ} are independent of τ.

[20] This simplifies the analysis without loss of generality as it avoids an explicit treatment of how indirect signals induce action recommendations.

2016)—for the case with a single Receiver who has no private information—to propose the notion of *Bayes correlated equilibrium*.[21]

However, unlike the existing information design literature, the current paper considers that (i) information design is restricted over any single separate aspect of uncertainty, and it does not affect the correlations between the aspects, and (ii) the Sender has private information about how the aspects correlate. Thus, building on the information design approach, let me define a *decision rule (under aspect-restricted commitment) over aspect $\kappa \in \{x, y\}$, when the Sender has type $\tau \in \mathcal{T}$ and chooses a message strategy $d \in \mathcal{D}$*, as a mapping $\sigma_\kappa^{(\tau,d)} : \mathcal{K} \to \Delta(A)$. The interpretation of a decision rule with aspect-restricted commitment $\sigma_\kappa^{(\tau,d)}$ is that if the true realization of the respective aspect κ is $\kappa_l \in \mathcal{K}$, then, contingent on making a decision $d \in \mathcal{D}$ on disclosure about his type, a Sender of type $\tau \in \mathcal{T}$ recommends action $a \in A$ with probability $\sigma_\kappa^{(\tau,d)}(a \mid \kappa_l)$.

A string $\sigma_\kappa^d \equiv (\sigma_\kappa^{(\tau,d)})_{\tau\in\mathcal{T}}$ gives us a generic *decision rule over aspect κ, contingent on the message strategy d*, and then $\sigma_\kappa = (\sigma_\kappa^d)_{d\in\mathcal{D}}$ indicates a possible *profile of decision rules over aspect κ*. Also, let me use $(\kappa; \sigma_\kappa)$ to denote an *investigation choice*, which identifies the aspect κ chosen for investigation, as well as a list of decision rules σ_κ^d, which are selected contingent on each possible message strategy $d \in \mathcal{D}$. Furthermore, since we are considering a two-action setting, in some parts of the paper it will be convenient for simplicity to use the short-hand notation $\hat{\sigma}_{k_l}^\tau \equiv \sigma_\kappa^{(\tau,r)}(\overline{a} \mid \kappa_l) \in [0,1]$ to denote, for the full disclosure strategy ($d = r$), the probability according to which investigation on aspect κ recommends acceptance of the proposal, given that the aspect realization is κ_l. Then, a generic decision rule over aspect κ, conditional on the Sender fully revealing his verifiable information about the pattern of dependence, can be described using a list of probabilities $\hat{\sigma}_\kappa^\tau \equiv \{\hat{\sigma}_{\kappa_l}^\tau \in [0,1]\}_{l=1}^m$. Likewise, I will sometimes use the short-hand notation $\hat{\sigma}_{k_l}^w \equiv \sigma_\kappa^{(\tau,w)}(\overline{a} \mid \kappa_l) \in [0,1]$ to denote, for the withholding strategy ($d = w$), the probability according to which investigation on aspect κ recommends acceptance of the proposal, given that the aspect realization is κ_l.

There is a clear analogy with the notion of decision rule under complete commitment. However, an aspect-restricted decision rule makes recommendations based only on partial information about the state. Compared to complete commitment, less information is disclosed under aspect-restricted commitment. Exactly as in the complete commitment benchmark, though, the approach rests on the consideration that the Sender does not need to know the true realization of the respective aspect κ. The commitment assumption crucially requires

[21] Equivalently, the requirements in Eq. (1) characterize the behavior of an information designer in the key concavification problem explored by Kamenica and Gentzkow (2011) for Bayesian persuasion.

that the Sender can condition the decision rule σ_κ^d on the realization of aspect κ.[22]

To fix ideas about how a decision rule under aspect-restricted commitment discloses credible information to the Receiver in the proposed setup, suppose that the Sender has type τ, makes an investigation choice $(\kappa; \sigma_\kappa)$, and chooses a message strategy d. The Receiver will use the Sender's message strategy d to form a system of beliefs β^d and, accordingly, to compute expectations $\mathbb{E}_{\beta_\tau^d}[\tilde{\tau}]$ about the pattern of dependence, for each actual type τ. Given these ingredients, incentive-compatibility, or obedience, requires then each profile of decision rules σ_κ to satisfy, for each type-message strategy pair $(\tau, d) \in \mathcal{T} \times \mathcal{D}$, the condition:

$$\sum_{\theta \in \Theta} \mathbb{E}_{\beta_\tau^d}\Big[\ \sigma_\kappa^{(\tilde{\tau},d)}(\overline{a} \mid \kappa)\, \psi_{\tilde{\tau}}(\theta)\ \Big] \big[u_R(\overline{a}, \theta) - u_R(\underline{a}, \theta)\big] \geq 0. \tag{2}$$

Notice that the condition in Eq. (2) above is simply an adjusted version, in terms of the expected value $\mathbb{E}_{\beta_\tau^d}[\cdot]$, of the key obedience condition in Eq. (1). In addition to requiring the computation of expected patterns of dependence $\mathbb{E}_{\beta_\tau^d}[\tilde{\tau}]$, the key difference is that the action recommendation in the proposed benchmark is based only on the realization of a particular dimension $\kappa \in \{x, y\}$ of the state θ.

Importantly, the Revelation Principle arguments provided by Bergemann and Morris (2016) in their proof of Proposition 1 apply entirely to the definition of decision rule $\sigma_\kappa^{(\tau,d)}$ under aspect-restricted commitment, for each given aspect κ and each given type and message strategy pair (τ, d). In particular, contrary to the insights of the literature that initiated with Bester and Strausz (2001), wherein the Revelation Principle is indeed challenged, the Sender is not limited in his commitment power over the chosen aspect. Therefore, the standard Revelation Principle does not fail in this case because the Sender cannot exploit the selected decision rule to his advantage as it is the case under *partial* or *limited* commitment power. The Receiver accordingly anticipates that any (possibly indirect) signal disclosed by the information structure corresponds to truthful reporting. Sender and Receiver commonly know that the recommendations from the selected information structures are completely binding, in spite of being based only on partial information about the state. Therefore, although we can consider in principle that investigation over a given aspect offers any (perhaps indirect, or through signal realizations) communication mechanism, we can further resort without loss of generality to a "direct communication" mechanism that recommends actions to the

[22] The premises behind the idea of a decision rule σ_κ^d can be intuitively phrased as: the two players commonly known that the Sender is both (i) able to "commission" any possible investigation process over aspect κ and (ii) unable to affect in any way the data subsequently released by the chosen investigation.

Receiver. Then, we only need to verify that the Receiver is given the right incentives to obey the recommendations from aspect-restricted investigation, as expressed in Eq. (2) above.

Notice that if the Sender reveals some information about his type, then he may select a different aspect for information design and/or provide a different information structure for each different type that he reveals to the Receiver. For instance, in the extreme case where $d = r$, if $\tau \neq \tau'$, the Sender can provide the Receiver with different decision rules, $\sigma_\kappa^{(\tau,r)}$ and $\sigma_{\kappa'}^{(\tau',r)}$. This is in total consonance with an informational consistency requirement where (ex ante) information design does not interfere with (interim) communication about the pattern of correlation. On the other hand, though, if the Sender chooses to withhold some of his verifiable private information about the pattern of correlation—i.e., $d(\tau) = \mathcal{T}_\tau \neq \{\tau\}$ for some non-singleton subset $\mathcal{T}_\tau \subset \mathcal{T}$—, then it must be the case that the Sender makes a common aspect choice, and provides a common information structure as well, for each possible type $\tau' \in d(\tau) = \mathcal{T}_\tau$. A natural informational consistency condition is needed in the proposed set up: whenever $d(\tau) = \mathcal{T}_\tau \neq \{\tau\}$ (for some non-singleton subset $\mathcal{T}_\tau \subset \mathcal{T}$), then it must be the case that $\sigma_\kappa^{(\tau',d)} = \sigma_\kappa^{(\tau'',d)}$ for each $\tau', \tau'' \in \mathcal{T}_\tau$ and each $\kappa \in \{x, y\}$. Without this condition, the Receiver would infer some information about the Sender's type by recognizing different investigation choices.[23]

2.6. Leading Example

Before presenting the specifics of the equilibrium notion used in the paper, this Subsection 2.6 illustrates the proposed benchmark with a specific problem in terms of the example, spelled out in the Introduction, of the automobile company where its consulting CEO wants to persuade the Board to launch their electric model into the new market.

There are nine possible states of the world $\theta = (x, y) \in \Theta = \{\theta_1, \ldots, \theta_9\} = X \times Y$, with $X = \{x_1, x_2, x_3\}$ and $Y = \{y_1, y_2, y_3\}$. Let $\theta_1 = (x_1, y_1)$, $\theta_2 = (x_1, y_2)$, $\theta_3 = (x_1, y_3)$, $\theta_4 = (x_2, y_1)$, $\theta_5 = (x_2, y_2)$, $\theta_6 = (x_2, y_3)$, $\theta_7 = (x_3, y_1)$, $\theta_8 = (x_3, y_2)$, and $\theta_9 = (x_3, y_3)$. Aspect x describes income growth and aspect y captures environmental concerns. The high action $\overline{a}$ is interpreted as launching the new model, whereas the low action $\underline{a}$ means rejecting such a proposal. The CEO wants to influence the Board's decision in favor of acceptance always, regardless of the true value of the state: $u_S(\overline{a}, \theta) = 1$ and $u_S(\underline{a}, \theta) = 0$ for each $\theta \in \Theta$. The Board wants to accept only if the state belongs to the acceptance set $\overline{\Theta} = \{\theta_3, \theta_6, \theta_8, \theta_9\}$.

[23] As a consequence, the Receiver will face a common incentive-compatibility constraint for each $\tau' \in d(\tau)$ with the form in Eq. (2), for an induced expected pattern of dependence $\mathbb{E}_{\beta^d_{\tau'}}[\tilde{\tau}]$, which in fact does not depend on the actual type $\tau' \in d(\tau)$ chosen by Nature.

In particular, consider that $u_R(\underline{a}, \theta) = 0$ for each $\theta \in \Theta$, whereas $u_R(\overline{a}, \theta) = 1/3$ if $\theta \in \overline{\Theta}$, and $u_R(\overline{a}, \theta) = -1$ if $\theta \in \Theta \backslash \overline{\Theta}$.[24] The marginal priors about the two aspects of the state are given by $\psi_x(x_1) = 3/6$, $\psi_x(x_2) = 1/6$ and $\psi_x(x_3) = 2/6$, and by $\psi_y(y_1) = \psi_y(y_2) = \psi_y(y_3) = 1/3$. The set of states, the acceptance set, and the marginal priors over the separate aspects for this example are shown in Fig. 1.

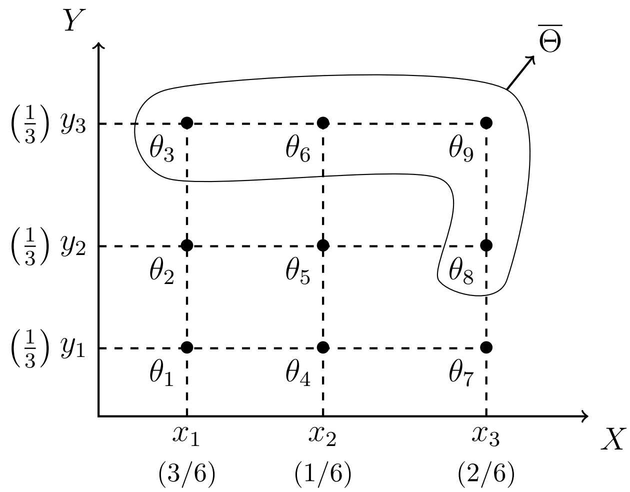

Figure 1 – Leading Example: Set of States and Marginal Priors.

There are two possible patterns of dependence, $\mathcal{T} = \{\tau_1, \tau_2\}$, that may relate aspects x and y, where

$$\tau_1 = \begin{bmatrix} 1/3 & 1/3 & 1/3 \\ 1/3 & 1/3 & 1/3 \\ 1/3 & 1/3 & 1/3 \end{bmatrix} \quad \text{and} \quad \tau_2 = \begin{bmatrix} 2/3 & 0 & 0 \\ 1/3 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

The prior q over possible patterns of dependence is given by $q(\tau_1) = 1/2$.

Note that both possible patterns τ_1 and τ_2 satisfy the key Bayesian plausibility condition of Assumption 1. The two aspects are independent under pattern τ_1, whereas there is some

[24] Note that preferences in this example are analogous versions (for a two-dimensional state of the world) of those in the leading "plaintiff-judge example" of Kamenica and Gentzkow (2011)'s influential contribution on (Bayesian) persuasion and "investment example" of Bergemann and Morris (2019)'s survey on information design.

degree of correlation between them under pattern τ_2. Also, since

$$\sum_{\theta\in\Theta}\psi_{\tau_1}(\theta)\ u_R(\overline{a},\theta) = (4/9)(1/3) + (5/9)(-1) < 0 \quad \text{and}$$
$$\sum_{\theta\in\Theta}\psi_{\tau_2}(\theta)\ u_R(\overline{a},\theta) = (1/3)(1/3) + (2/3)(-1) < 0,$$

it follows that $\hat{a}_R(\tau) = \underline{a}$ for each possible pattern of dependence τ_1 and τ_2. Using any possible priors ψ_τ about θ, the Board always wants to reject. The details of this example satisfy all the assumptions of the proposed benchmark, including Assumption 1 and Assumption 2.

Let us study how the CEO would optimally decide about investigation over the separate aspects x and y, and communicate about the pattern of dependence τ. In this example, we just need to study the message strategies $d \in \{w, r\}$, where the CEO either pools completely or reveals fully his private information about the prior ψ_τ. Conditional on the CEO's message strategy $d \in \{w, r\}$, consider a pair of beliefs β^d_τ, for each $\tau \in \{\tau_1, \tau_2\}$, which we can parameterize as $\beta^d_\tau(\tau_1) = 1 - \varepsilon$ for some $\varepsilon \in [0, 1]$. Thus, ε gives us the probability of facing the prior that features some degree of correlation between the two aspects. It follows that

$$\mathbb{E}_{\beta^d_\tau}[\tilde{\tau}] = (1/3)\begin{bmatrix} 1+\varepsilon & 1-\varepsilon & 1-\varepsilon \\ 1 & 1+2\varepsilon & 1-\varepsilon \\ 1-\varepsilon & 1-\varepsilon & 1+2\varepsilon \end{bmatrix}.$$

Observe that $d = w$ would *in principle*[25] lead to that the Board retains her priors q and, therefore, to that $\varepsilon = 1/2$. On the other hand, $d = r$ leads to $\varepsilon = 0$ when $\tau = \tau_1$, and to $\varepsilon = 1$ when $\tau = \tau_1$. The expected prior $\mathbb{E}_{\beta^d_\tau}[\psi_{\tilde{\tau}}(\theta)]$ about the state θ can then be computed as:

$$\begin{aligned} &\mathbb{E}_{\beta^d_\tau}[\psi_{\tilde{\tau}}(\theta_1)] = 3(1+\varepsilon)/18, \quad \mathbb{E}_{\beta^d_\tau}[\psi_{\tilde{\tau}}(\theta_2)] = 3/18, \quad \mathbb{E}_{\beta^d_\tau}[\psi_{\tilde{\tau}}(\theta_3)] = 3(1-\varepsilon)/18; \\ &\mathbb{E}_{\beta^d_\tau}[\psi_{\tilde{\tau}}(\theta_4)] = \mathbb{E}_{\beta^d_\tau}[\psi_{\tilde{\tau}}(\theta_6)] = (1-\varepsilon)/18, \quad \mathbb{E}_{\beta^d_\tau}[\psi_{\tilde{\tau}}(\theta_5)] = (1+2\varepsilon)/18; \\ &\mathbb{E}_{\beta^d_\tau}[\psi_{\tilde{\tau}}(\theta_7)] = \mathbb{E}_{\beta^d_\tau}[\psi_{\tilde{\tau}}(\theta_8)] = 2(1-\varepsilon)/18, \quad \mathbb{E}_{\beta^d_\tau}[\psi_{\tilde{\tau}}(\theta_9)] = 2(1+2\varepsilon)/18. \end{aligned} \tag{3}$$

Following the expression derived in Eq. (3) above, the possible priors about the state θ that correspond, respectively, to the patterns of dependence τ_1 and τ_2 are depicted in Fig. 2. Likewise, the expected prior about the state θ induced by the prior q over the possible

[25] As we will analyze later, an skeptical Receiver could infer something totally different when she learns that the informed Sender withholds his private information.

patterns of dependence, is depicted in Fig. 3.

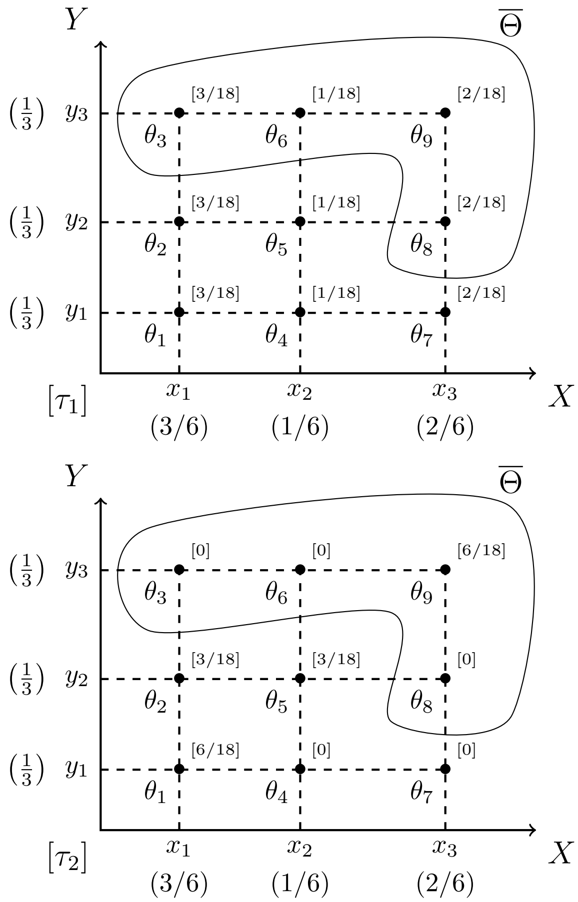

Figure 2 – Leading Example: Possible Priors about θ.

For the case of full revelation of the CEO's private information ($d = r$), investigation over each aspect $\kappa \in \{x, y\}$ can be described by setting three parameters for each of the types $\tau \in \{\tau_1, \tau_2\}$. As suggested earlier in Subsection 2.5, for $j, k = 1, 2, 3$, let me use the short-hand notations $\hat{\sigma}^{\tau}_{x_j} \equiv \sigma^{(\tau,r)}_x(\overline{a} \mid x_j) \in [0, 1]$ and $\hat{\sigma}^{\tau}_{y_k} \equiv \sigma^{(\tau,r)}_y(\overline{a} \mid y_k) \in [0, 1]$. Likewise, for the case of withholding ($d = w$), investigation can also be described, as suggested earlier, by setting three parameters that are independent of the CEO's type: $\hat{\sigma}^{w}_{x_j} \equiv \sigma^{(\tau,w)}_x(\overline{a} \mid x_j) \in [0, 1]$ and $\hat{\sigma}^{w}_{y_k} \equiv \sigma^{(\tau,w)}_y(\overline{a} \mid y_k) \in [0, 1]$.

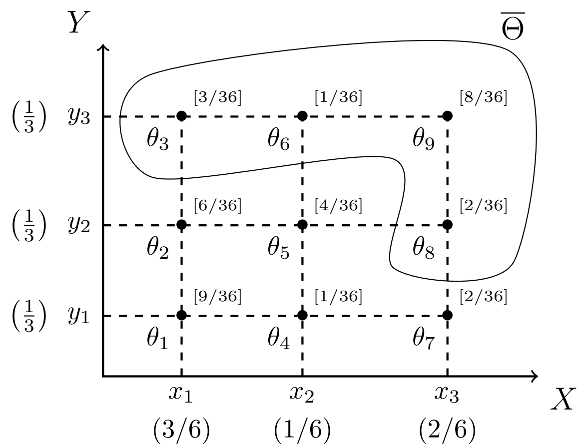

Figure 3 – Leading Example: Expected Prior about θ under q.

2.6.1. Investigation over Income Growth

Take an investigation choice $(x; \sigma_x)$. The incentive-compatibility condition detailed earlier in Eq. (2) becomes:

$$\begin{aligned} &\left[24\mathbb{E}_{\beta^d_\tau}[\psi_{\tilde{\tau}}(\theta_3)] - 9\right]\sigma_x^{(\tau,d)}(\overline{a} \mid x_1) \\ &+ \left[24\mathbb{E}_{\beta^d_\tau}[\psi_{\tilde{\tau}}(\theta_6)] - 3\right]\sigma_x^{(\tau,d)}(\overline{a} \mid x_2) + \left[2 - 24\mathbb{E}_{\beta^d_\tau}[\psi_{\tilde{\tau}}(\theta_7)]\right]\sigma_x^{(\tau,d)}(\overline{a} \mid x_3) \geq 0. \end{aligned} \tag{4}$$

In addition, observe that the (ex ante) expected utility that the CEO receives, when he is type τ and chooses message strategy d, is given by:

$$\sigma_x^{(\tau,d)}(\overline{a} \mid x_1)(3/6) + \sigma_x^{(\tau,d)}(\overline{a} \mid x_2)(1/6) + \sigma_x^{(\tau,d)}(\overline{a} \mid x_3)(2/6).$$

Consider first the situation where that the CEO chooses $d = r$, so that he fully reveals his type and the Board learns the true value of the pattern of dependence τ. Using the expected prior beliefs $\mathbb{E}_{\beta^d_\tau}[\psi_{\tilde{\tau}}(\theta)]$ derived in Eq. (3), the incentive-compatibility requirement on information design in Eq. (4) above turns into:

- For $\tau = \tau_1$ (by considering $\varepsilon = 0$),

$$15\hat{\sigma}^{\tau_1}_{x_1} + 5\hat{\sigma}^{\tau_1}_{x_2} + 2\hat{\sigma}^{\tau_1}_{x_3} \leq 0.$$

  In this case, all that the CEO can do when trying to maximize his expected utility is to select $\hat{\sigma}^{\tau_1}_{x_1} = \hat{\sigma}^{\tau_1}_{x_2} = \hat{\sigma}^{\tau_1}_{x_3} = 0$. Thus, when both aspects are independent and the

Boards learns so, investigation over income is unable to influence the Board towards acceptance and the CEO obtains a zero expected utility.

- For $\tau = \tau_2$ (by considering $\varepsilon = 1$),

$$2\hat{\sigma}^{\tau_2}_{x_3} \geq 9\hat{\sigma}^{\tau_2}_{x_1} + 3\hat{\sigma}^{\tau_2}_{x_2}.$$

  In this case, the CEO would maximize his expected utility by selecting $\hat{\sigma}^{\tau_2}_{x_3} = 1$, and any $\hat{\sigma}^{\tau_2}_{x_1}, \hat{\sigma}^{\tau_2}_{x_2} \in [0,1]$ such that $(9/2)\hat{\sigma}^{\tau_2}_{x_1} + (3/2)\hat{\sigma}^{\tau_2}_{x_2} = 1$. His expected utility would then be 4/9.

Consider now the situation that the CEO chooses $d = w$, so that he withholds his information about τ and the Board would in principle retain her priors q. From the expression of the expected priors $\mathbb{E}_{\beta^d_\tau}[\psi_{\tilde{\tau}}(\theta)]$ derived in Eq. (3), the incentive-compatibility condition on information design in Eq. (4) above turns into:

$$2\hat{\sigma}^{w}_{x_3} \geq 21\hat{\sigma}^{w}_{x_1} + 7\hat{\sigma}^{w}_{x_2}.$$

In this case, the CEO would maximize his expected utility by selecting $\hat{\sigma}^{w}_{x_3} = 1$, and any $\hat{\sigma}^{w}_{x_1}, \hat{\sigma}^{w}_{x_2} \in [0,1]$ such that $(21/2)\hat{\sigma}^{w}_{x_1} + (7/2)\hat{\sigma}^{w}_{x_2} = 1$. For each possible type $\tau \in \{\tau_1, \tau_2\}$, the CEO's expected utility would then be

$$\hat{\sigma}^{w}_{x_1}(3/6) + \hat{\sigma}^{w}_{x_2}(1/6) + \hat{\sigma}^{w}_{x_3}(2/6) = 8/21.$$

In order to describe how the players behave in equilibrium, we still need to study the beliefs over the pattern of dependence that a skeptical Board forms in the event that the CEO decides to not reveal his information. A skeptical Board places herself in the CEO's position and computes her own expected payoffs induced by each possible type of the CEO, conditional on discounting the optimal investigation choice followed by the CEO for each of his types. Then, skepticism makes the Board place positive probability only on the type(s) that give(s) her the lowest expected utility. The expected utility of the Board derived for each possible type of the CEO, given that the CEO follows his optimal investigation choice for the respective type, is:

- For $\tau = \tau_1$, the Board receives a zero expected utility since the CEO is not able to persuade towards acceptance for the low type.

- For $\tau = \tau_2$, the Board receives the expected utility

$$\hat{\sigma}^{\tau_2}_{x_1}(3/6)[-1] + \hat{\sigma}^{\tau_2}_{x_2}(1/6)[-1] + \hat{\sigma}^{\tau_2}_{x_3}(2/6)[1/3] = 0,$$

since the CEO can in this case optimally select, for example, $\hat{\sigma}^{\tau_2}_{x_1} = 0$, $\hat{\sigma}^{\tau_2}_{x_2} = 2/3$, and $\hat{\sigma}^{\tau_2}_{x_3} = 1$.

Since investigation over income growth leaves the skeptical Board indifferent between the two possible types of the CEO, she then believes that *any* possible type could be the actual one. In particular, revising her beliefs according to her priors and, therefore, considering $\varepsilon = 1/2$ is the most neutral position which is compatible with such a skeptical Bayesian updating. Then, given this updating, since the type τ_2 of the CEO receives a higher expected utility from verifiably disclosing itself (4/9), compared to what the obtains by pooling alongside with type τ_1 (8/21), the CEO has strict incentives to choose $d = r$ and reveal completely his private information about how income growth and environmental concerns are related. Observe that this will be the case regardless of whether communication is pure cheap talk or through verifiable messages.

2.6.2. Investigation over Environmental Concerns

Now I turn to explore what can be attained through investigation over aspect y. Take an investigation choice $(y; \sigma_y)$. The incentive-compatibility condition detailed earlier in Eq. (2) now becomes:

$$-3\sigma^{(\tau,d)}_y(\overline{a} \mid y_1) + \left[12\mathbb{E}_{\beta^d_\tau}[\psi_{\tilde{\tau}}(\theta_8)] - 3\right]\sigma^{(\tau,d)}_y(\overline{a} \mid y_2) + \sigma^{(\tau,d)}_y(\overline{a} \mid y_3) \geq 0. \qquad (5)$$

In addition, observe that the (ex ante) expected utility that the CEO receives, when he is type τ and chooses message strategy d, is given by:

$$\sigma^{(\tau,d)}_y(\overline{a} \mid y_1)(1/3) + \sigma^{(\tau,d)}_y(\overline{a} \mid y_2)(1/3) + \sigma^{(\tau,d)}_y(\overline{a} \mid y_3)(1/3).$$

Suppose first that the CEO chooses $d = r$. Using the expected priors $\mathbb{E}_{\beta^d_\tau}[\psi_{\tilde{\tau}}(\theta)]$ derived in Eq. (3), the incentive-compatibility requirement on information design in Eq. (5) above turns into:

- For $\tau = \tau_1$ (by considering $\varepsilon = 0$),

$$\hat{\sigma}^{\tau_1}_{y_3} \geq 3\hat{\sigma}^{\tau_1}_{y_1} + (5/3)\hat{\sigma}^{\tau_1}_{y_2}.$$

In this case, the CEO would maximize his expected utility by selecting $\hat{\sigma}^{\tau_1}_{y_3} = 1$, $\hat{\sigma}^{\tau_1}_{y_1} = 0$, and $\hat{\sigma}^{\tau_1}_{y_2} = 3/5$. His ex ante expected utility would then be 8/15.

- For $\tau = \tau_2$ (by considering $\varepsilon = 1$),

$$\hat{\sigma}^{\tau_2}_{y_3} \geq 3\hat{\sigma}^{\tau_2}_{y_1} + 3\hat{\sigma}^{\tau_2}_{y_2}.$$

  In this case, the CEO would maximize his expected utility by selecting $\hat{\sigma}^{\tau_2}_{y_3} = 1$, and any $\hat{\sigma}^{\tau_2}_{y_1}, \hat{\sigma}^{\tau_2}_{y_2} \in [0, 1]$ such that $3\hat{\sigma}^{\tau_2}_{y_1} + 3\hat{\sigma}^{\tau_2}_{y_2} = 1$. His ex ante expected utility would be 4/9.

Suppose now that the CEO chooses $d = w$. We observe from the expected priors $\mathbb{E}_{\beta^d_{\tau}}[\psi_{\tilde{\tau}}(\theta)]$ derived in Eq. (3) that the required incentive-compatibility condition on information design in Eq. (5) above turns into:

$$\hat{\sigma}^{w}_{y_3} \geq 3\hat{\sigma}^{w}_{y_1} + (7/3)\hat{\sigma}^{w}_{y_2}.$$

In this case, the CEO would maximize his expected utility by selecting $\hat{\sigma}^{w}_{y_3} = 1$, $\hat{\sigma}^{w}_{y_1} = 0$, and $\hat{\sigma}^{w}_{y_2} = 3/7$. For each possible type $\tau \in \{\tau_1, \tau_2\}$, the CEO's expected utility would then be

$$\hat{\sigma}^{w}_{y_1}(1/3) + \hat{\sigma}^{w}_{y_2}(1/3) + \hat{\sigma}^{w}_{y_3}(1/3) = 10/21.$$

Again, we need to study the beliefs over the pattern of dependence that a skeptical Board forms in the event that the CEO decides to not disclose his information. The expected utility of the Board derived for each possible type of the CEO, given that the CEO follows his optimal investigation choice for the respective type, is:

- For $\tau = \tau_1$, the Board receives

$$\hat{\sigma}^{\tau_1}_{y_1}(1/3)[-1] + \hat{\sigma}^{\tau_1}_{y_2}(1/3)\{(4/6)[-1] + (2/6)[1/3]\} + \hat{\sigma}^{\tau_1}_{y_3}(1/3)[1/3] = 0,$$

  since the CEO is optimally selecting in this case $\hat{\sigma}^{\tau_1}_{y_1} = 0$, $\hat{\sigma}^{\tau_1}_{y_2} = 3/5$, and $\hat{\sigma}^{\tau_1}_{y_3} = 1$.

- For $\tau = \tau_2$, the Board receives

$$\hat{\sigma}^{\tau_2}_{y_1}(3/6)(2/3)[-1] + \hat{\sigma}^{\tau_2}_{y_2}\{(3/6)(1/3)[-1] + (1/6)[-1]\} + \hat{\sigma}^{\tau_2}_{y_3}(2/6)[1/3] = 0,$$

since the CEO can in this case optimally select, for example, $\hat{\sigma}_{y_1}^{\tau_2} = 0$, $\hat{\sigma}^{\tau_2} = 1/3$, and $\hat{\sigma}_{y_3}^{\tau_2} = 1$.

Investigation over environmental concerns makes the skeptical Board indifferent between the possible patterns of dependence as well. Therefore, the Board believes that any possible type is the true type of the CEO. This is a general insight from the proposed setup because, under the maintained assumptions, optimal (aspect-restricted) information design disciplines the Receiver in a way such that she receives a constant utility across the various possible types of the Sender. This implication follows from the fact that the information design problem of the Sender has a unique incentive-compatibility constraint in the proposed two-action (either accept or reject) setting. In particular, updating according to her priors and, therefore, considering $\varepsilon = 1/2$ is compatible with such a skeptical Bayesian updating and simply leaves the CEO with her priors about τ. Given this, since the type τ_1 of the CEO receives a higher expected utility from verifiably disclosing itself (8/15), compared to what the obtains by pooling alongside with type τ_2 (10/21), the CEO has strict incentives to choose $d = r$ and reveal completely his private information about how income and environmental concerns are related. Again, this will be the case regardless of whether the reports can be verified or not.

We observe that revealing his private information about τ gives a CEO of type τ_1 an expected utility of zero, when he commissions investigation over income, and an expected utility of 8/15, when he commissions investigation over environmental concerns. Therefore, if $\tau = \tau_1$ (so that the two aspects of uncertainty are independent), the CEO optimally chooses environmental concerns for information design. On the other hand, when he is type τ_2, and reveals such a private piece of information, the CEO receives an expected utility of 4/9, regardless of whether he chooses investigation over income growth or over environmental concerns. Therefore, if $\tau = \tau_2$ (so that the two aspects of uncertainty are correlated), the CEO is indifferent between choosing income growth or environmental concerns for information design.

To summarize, the equilibrium disclosure behavior by the CEO in this example involves:

(i) selecting investigation over environmental concerns when the two aspects are unrelated, and over any of the two aspects when there is some correlation between them;

(ii) for the selected aspect, choosing an information structure such that:

(iia) for the pattern of correlation τ_1, investigation over environmental concerns always recommends rejection when $y = y_1$, recommends acceptance with probability 3/5 when $y = y_2$, and always recommends acceptance when $y = y_3$, whereas,

(iib) for the pattern of correlation τ_2, investigation over environmental concerns always recommends rejection when $y = y_1$, recommends acceptance with probability 1/3 when $y = y_2$, and always recommends acceptance when $y = y_3$, and investigation over income growth always recommends acceptance when $x = x_3$, and randomize for the cases where $x = x_1$ and $x = x_2$ (in a way such that the overall probability of acceptance the probability conditional on either $x = x_1$, $x = x_2$, or both equals 1/3); and

(iii) fully revealing his private information about how environmental concerns and income are related in the targeted market.

3. Equilibrium

The equilibrium notion can perhaps be best understood from the Sender's perspective. The Sender's disclosure behavior has two effects. First, information design places incentive-compatibility conditions, in terms of additional information, for the Receiver to follow the recommendations of investigation over the selected aspect. Secondly, (interim) communication about the Sender's type leads to posteriors about the pattern of dependence, based on a skeptical Bayesian belief updating process followed by the Receiver. The combination of both sources of new information determines the Receiver's optimal action. The suitable solution concept in this environment must therefore add incentive-compatibility requirements to the usual conditions that specify a *(weak) perfect Bayes-Nash equilibrium*.

Let me use a list $\phi = ((\sigma_x, \sigma_y); d, \kappa)$ to identify an *information disclosure strategy*, which consists of (i) profiles of decision rules σ_κ for each of the two possible aspects $\kappa \in \{x, y\}$, (ii) a message strategy $d \in \mathcal{D}$ about the relationships τ between the two aspects, and (iii) an aspect choice $\kappa \in \{x, y\}$ for investigation. Suppose that the Sender makes an investigation choice $(\kappa; \sigma_\kappa)$ and selects a message strategy d that induces a system of posteriors β^d. Then, for each type $\tau \in \mathcal{T}$ and each aspect $\kappa \in \{x, y\}$, the Sender's information design problem needs to satisfy the incentive-compatibility, or obedience, condition derived earlier in Eq. (2). Under such a restriction, the Sender will then want to select decision rules that maximize his expected utility. Conditional on the type of the Sender being τ, let $U_S(\phi \mid \tau)$ be his expected utility for an information disclosure strategy ϕ, which is given by

$$U_S(\phi \mid \tau) \equiv \sum_{a \in A} \sum_{\theta \in \Theta} \sigma_\kappa^{(\tau,d)}(a \mid \kappa)\ \psi_\tau(\theta)\ u_S(a,\theta). \tag{6}$$

On the other hand, conditional on the type of the Sender being τ, let $U_R(\phi \mid \tau)$ be the Receiver's expected utility under an information disclosure strategy ϕ, which is given by[26]

$$U_R(\phi \mid \tau) \equiv \sum_{a \in A} \sum_{\theta \in \Theta} \mathbb{E}_{\beta_\tau^d}\big[\,\sigma_\kappa^{(\tilde{\tau},d)}(a \mid \kappa)\ \ \psi_{\tilde{\tau}}(\theta)\,\big]\ u_R(a,\theta). \tag{7}$$

The model considers that the Receiver's updating rule is Bayesian and, in addition, that it follows a maximally skeptical—or precautionary—approach.[27] In particular, for the case where the Sender pools over a subset of his possible types, the Receiver places positive probability only on the Sender's type(s) that lead(s) herself to the lowest possible expected payoffs, conditional on the optimal investigation choice followed by the Sender for each possible type. Moreover, for those cases where several distinct types meet this skepticism criterion, the Receiver simply updates her beliefs by assigning each type a proportional probability, according to her priors about types, among the possible types under the newly disclosed information. Of course, this is not the unique updating rule that could be followed when the Receiver is indifferent between several types of the Sender but it seems to be the most neutral approach that incorporates both Bayesian updating and a maximally skeptical position by the Receiver. Algebraically,

Assumption 3. *If the Sender selects an information disclosure strategy $\phi = ((\sigma_x, \sigma_y); d, \kappa)$ such that $d \neq r$, then the Receiver updates her prior beliefs q about the pattern of dependence by selecting, for each type $\tau \in \mathcal{T}$ such that $d(\tau) = \mathcal{T}_\tau \neq \{\tau\}$ (for some non-singleton subset $\mathcal{T}_\tau \subset \mathcal{T}$) and for each type $\tau' \in \mathcal{T}_\tau$, the particular posterior beliefs $\beta_{\tau'}^d$ such that $\beta_{\tau'}^d(\tau^*) > 0$ only for types $\tau^* \in \mathcal{T}$ that satisfy*

$$\tau^* \in \mathcal{T}_{\text{skep}}(\tau) \equiv arg \min_{\{\tau' \in \mathcal{T}_\tau\}} \quad \sum_{a \in A} \sum_{\theta \in \Theta} \sigma_\kappa^{(\tau',r)^*}(a \mid \kappa)\ \psi_{\tau'}(\theta)\ u_R(a,\theta), \tag{8}$$

[26] Of course, since the model assumes that $u_S(\underline{a},\theta) = u_R(\underline{a},\theta) = 0$ for each $\theta \in \Theta$, the expressions in Eq. (6) and Eq. (7) can be respectively simplified to $U_S(\phi \mid \tau) = \sum_{\theta \in \Theta} \sigma_\kappa^{(\tau,d)}(\overline{a} \mid \kappa)\ \psi_\tau(\theta)\ u_S(\overline{a},\theta)$ and $U_R(\phi \mid \tau) = \sum_{\theta \in \Theta} \mathbb{E}_{\beta_\tau^d}\big[\,\sigma_\kappa^{(\tilde{\tau},d)}(\overline{a} \mid \kappa)\ \ \psi_{\tilde{\tau}}(\theta)\,\big]\ u_R(\overline{a},\theta)$. I presented the general expression in the main text, though, to introduce the general form of the players' expected utilities in the information design problem under more general preferences.

[27] As in the classical literature on persuasion (Grossman, 1981; Milgrom, 1981; Milgrom and Roberts, 1986).

conditional on the respective decision rule $\sigma_\kappa^{(\tau',r)^*}$ being incentive-compatible constrained optimal for the Sender when he has type τ' and fully reveals it. Moreover, if some set $\mathcal{T}_{\text{skep}}(\tau)$ is not a singleton, then the Receiver assigns probability $\beta_\tau^d(\tau') = q(\tau')/\sum_{\tau''\in\mathcal{T}_{\text{skep}}(\tau)} q(\tau'')$ to each possible type $\tau' \in \mathcal{T}_{\text{skep}}(\tau)$.

In Assumption 3, I am considering that the Receiver computes her possible expected utilities, for each type, provided that she actually learns the true type. This is why we need to take the optimal decision rule $\sigma_\kappa^{(\tau',r)^*}$ under full revelation (i.e., for $d = r$) for each $\tau' \in d(\tau) = \mathcal{T}_\tau \neq \{\tau\}$ (for some non-singleton subset $\mathcal{T}_\tau \subset \mathcal{T}$).

The overall persuasion goal of the Sender involves then selecting first a decision rule $\sigma_\kappa^{(\tau,d)}$ (for each type $\tau \in \mathcal{T}$, for each message strategy $d \in \mathcal{D}$, and for each aspect $\kappa \in \{x, y\}$) in a way such that the expected utility $U_S(\phi \mid \tau)$ derived in Eq. (6) is maximized under the constraint given in Eq. (2). Secondly, the description of the Sender's optimal disclosure behavior is closed by requiring him to select, for each possible type, both the message strategy d and the aspect κ that yield him the higher (ex ante) expected utility.

The description of the Sender's optimal behavior, and of the Receiver's updating rules and best reply, spelled out above—and illustrated in our leading example in Subsection 2.6—is formally stated in Definition 1 below.

Definition 1. An *equilibrium* of the described information disclosure game is an information disclosure strategy $\phi^* = ((\sigma_x^*, \sigma_y^*); d^*, \kappa^*)$ and a collection of posterior beliefs β^{d^*} such that:

(i) *(Ex ante) optimality constrained to incentive-compatibility, or obedience*: for each possible type $\tau \in \mathcal{T}$, for each aspect $\kappa \in \{x, y\}$, and each disclosure choice $d \in \mathcal{D}$, the decision rule $\sigma_\kappa^{(\tau,d)^*}$ solves the problem

$$\begin{aligned} \max_{\{\sigma_\kappa^{(\tau,d)}\}} \quad & \sum_{a\in A}\sum_{\theta\in\Theta} \sigma_\kappa^{(\tau,d)}(a \mid \kappa)\ \psi_\tau(\theta)\ u_S(a,\theta) \\ \text{s.t.:} \quad & \mathbb{E}_{\beta_\tau^d}\Big[\sum_{\theta\in\Theta} \sigma_\kappa^{(\tilde{\tau},d)}(\overline{a} \mid \kappa)\ \psi_{\tilde{\tau}}(\theta)\Big]\ \big[u_R(\overline{a},\theta) - u_R(\underline{a},\theta)\big] \geq 0. \end{aligned} \tag{9}$$

(ii) *Consistency of posterior beliefs about patterns of dependence:* upon $d^* \neq r$ the posterior beliefs $\beta_\tau^{d^*}$ (for each type $\tau \in \mathcal{T}$) are derived from q and d^* according to Bayes' rule and, in addition, following the skepticism condition given by Assumption 3.

(iii) *(Interim) optimality of message strategy*: there is no other $d' \in \mathcal{D}$ such that, for each

$\tau \in \mathcal{T}$,

$$U_S(((\sigma_x^*, \sigma_y^*); d', \kappa^*) \mid \tau) > U_S(((\sigma_x^*, \sigma_y^*); d^*, \kappa^*) \mid \tau).$$

(iv) *Optimal aspect choice*: type $\tau \in \mathcal{T}$ of the Sender selects aspect κ^* for investigation whenever

$$U_S(((\sigma_x^*, \sigma_y^*); d^*, \kappa^*) \mid \tau) \geq U_S(((\sigma_x^*, \sigma_y^*); d^*, -\kappa) \mid \tau).$$

Equilibria in the proposed information disclosure game will typically be multiple. This is an unsurprising implication when an information designer chooses decision rules in the classical information design framework. The optimization problem of the Sender has the form of a linear programming problem and, as usual in this class of problems, one can obtain multiple decision rules that satisfy the required incentive-compatibility condition and, at the same time, provide the Sender with a common optimal expected utility. In geometric terms, the usual tangency between the Sender's indifference line and the incentive-compatibility constraint gives this result. The assumption of information design restricted to a single dimension of uncertainty does not change the linear structure of the Sender's problem and, in a way totally analogous to the traditional information design benchmark, it allows for multiple aspect-restricted decision rules that solve his information design problem. In fact, our leading example in Subsection 2.6 featured this kind of multiplicity.

On the other hand, although strategic communication situations typically suffer from drastic forms of equilibria multiplicity, this element of the model will not contribute, under the maintained assumptions, to equilibria multiplicity in the proposed game. In particular, as Theorem 2 shows, the combination of information design over a single aspect with the way in which the Receiver skeptically updates her priors about the pattern of dependence lead to that each equilibrium features full disclosure of the pattern of dependence known by the Sender.

3.1. Aspect-Restricted *versus* Complete Information Design

A couple of comments might be helpful to appreciate better what the proposed setting allows the Sender to attain, relative to the traditional framework. First, notice that we would be able to suitably compare what decision rules under aspect-restricted commitment can attain relative to the complete commitment benchmark only for the case where the Receiver has full information about the pattern of dependence τ. Crucially, if the pattern of dependence τ is unknown by the Receiver, then the traditional notion of decision rule (under complete commitment) does not aptly capture information disclosure in the proposed benchmark if we wish to preserve the assumption that the Receiver is uncertain about the pattern of dependence. A decision rule under complete commitment will necessarily provide full information about the pattern of dependence τ as well since it is based on the joint realization $\theta = (x, y)$. For this reason, the benchmark proposed in this paper does not have a direct counterpart in the traditional information design setup.

Secondly, for the case where comparisons can indeed be suitably made—i.e., the Sender chooses $d = r$ so that the Receiver learns the true realization of his type τ—, any contingent action recommendation that can be achieved by a decision rule $\sigma_\kappa^{(\tau,r)}$ over any aspect κ, for any pattern of dependence τ, can also be achieved by a decision rule under complete commitment ρ^τ. This is very intuitive. All that a decision rule ρ^τ needs to do in order to offer the same contingent action recommendation as a decision rule with aspect-restricted commitment $\sigma_\kappa^{(\tau,r)}$ is simply to not condition its recommendations on the remaining random variable $-\kappa$. An information designer with the ability to design information over the entire state of the world $\theta = (x, y)$ has more flexibility—or, equivalently, is less (incentive-compatible) constrained—and, therefore, can attain higher ex ante utility. The formal arguments are detailed in Observation 1 below.

Observation 1. *Consider first a Sender that chooses investigation in the aspect-restricted setting proposed in this paper. Suppose, without loss of generality, that such a Sender (i) selects aspect $\kappa = x$ over which to design information, (ii) decides to fully disclose his type ($d = r$), and (iii) (for his given type τ) selects a decision rule $\sigma_x^{(\tau,r)^*}$ to maximize his ex ante expected utility (subject to the aspect-restricted incentive-compatibility constraint):*

$$\begin{aligned} &\max_{\{\sigma_x^{(\tau,r)}\}} \sum_{a\in A} \sum_{(x,y)\in\Theta} \sigma_x^{(\tau,r)}(a \mid x)\ t_{xy}\ \psi_x(x)\ u_S(a,(x,y)) \\ &s.t.: \sum_{(x,y)\in\Theta} \sigma_x^{(\tau,r)}(\overline{a} \mid x)\ t_{xy}\ \psi_x(x)\ \left[u_R(\overline{a},(x,y)) - u_R(\underline{a},(x,y))\right] = 0. \end{aligned} \tag{10}$$

The incentive-compatibility condition in problem Eq. (10) above is derived from Eq. (2) by making use of the equivalence $\psi_\tau((x, y)) = t_{xy}\ \psi_x(x)$, by considering that $d = r$, and by noting that such a restriction must hold with equality at the optimal information design choice of the Sender since he faces a single constraint to his maximization problem.

Secondly, instead of dealing with the aspect-restricted investigation setting, consider now a Sender with the ability to design information conditional on the the entire state of the world. The problem of such a Sender consists of choosing decision rules (under complete com-

mitment) ρ^τ *that satisfy the incentive-compatibility condition expressed earlier in Eq. (1). Notably, for a probabilistic environment where patterns of dependence satisfy the Bayesian plausibility requirement in Assumption 1, the Radon-Nikodym Theorem allows us to apply the definition of conditional probability to construct a decision rule under complete commitment* ρ^τ *from any given a decision rule under aspect-restricted commitment* $\sigma_x^{(\tau,r)}$ *by suitably selecting a family of conditional probability distributions* $\delta^\tau \equiv \{\delta^\tau(\cdot \mid \overline{a}, x) \in \Delta(Y) \mid x \in X\}$ *so as to satisfy the following Bayesian plausibility condition:*

$$\rho^\tau(\overline{a} \mid \theta) = (1/t_{xy})\ \delta^\tau(y \mid \overline{a}, x)\ \sigma_x^{(\tau,r)}(\overline{a} \mid x) \qquad \forall\theta = (x, y) \in \Theta. \tag{11}$$

Therefore, taking as given a Sender's optimal decision rule $\sigma_x^{(\tau,r)^*}$ *under aspect-restricted investigation, the incentive-compatibility constraint that a Sender with complete commitment power faces can be rewritten as:*

$$\sum_{\theta\in\Theta} \delta^\tau(y \mid \overline{a}, x)\ \sigma_x^{(\tau,r)^*}(\overline{a} \mid x)\ (1/t_{xy})\ \psi_\tau(\theta)\ \big[u_R(\overline{a}, \theta) - u_R(\underline{a}, \theta)\big] = 0.$$

In other words, a Sender with the ability to choose investigation simultaneously informative about both dimensions of uncertainty must select a decision rule $\rho^{\tau*}$ *which, crucially, can be constructed by picking a decision rule* $\sigma_x^{(\tau,r)^*}$ *and a family of conditional distributions* $\delta^{\tau*}$ *that satisfy the Bayesian condition requirement in Eq. (11) above. In addition, such families of conditional distributions* $\sigma_x^{(\tau,r)^*}$ *and* $\delta^{\tau*}$ *must be chosen in order to solve the problem:*

$$\begin{aligned} &\max_{\{\sigma_x^{(\tau,r)},\ \delta^\tau\}} \sum_{a\in A} \sum_{(x,y)\in\Theta} \sigma_x^{(\tau,r)}(a \mid x)\ \delta^\tau(y \mid \overline{a}, x)\ \psi_x(x)\ u_S(a, (x, y)) \\ &s.t.: \quad \sum_{(x,y)\in\Theta} \sigma_x^{(\tau,r)}(\overline{a} \mid x)\ \delta^\tau(y \mid \overline{a}, x)\ \psi_x(x)\ \big[u_R(\overline{a}, (x, y)) - u_R(\underline{a}, (x, y))\big] = 0. \end{aligned} \tag{12}$$

Although, as noted earlier, the two frameworks are not comparable in general, we observe that a Sender under the traditional information design has more flexibility relative to an information designer in the proposed approach of aspect-restricted commitment. In particular, upon selecting $\delta^\tau(y \mid \overline{a}, x) = t_{xy}$*, and then choosing* $\sigma_x^{(\tau,r)^*}$*, a Sender of type* τ *in the complete commitment environment has the ability to solve the problem in Eq. (10).*[28] *Conversely, a Sender in the aspect-restricted world cannot solve the full commitment problem in Eq. (12)*

[28] Importantly, by choosing $\delta^\tau(y \mid \overline{a}, x) = t_{xy}$, the Sender will not necessarily solve the information design problem under full commitment (in Eq. (12)) but he will certainly solve the corresponding problem under aspect-restricted commitment (in Eq. (10)).

simply because he does not have the ability to choose the family of distributions δ^{τ}.

To illustrate these insights, Subsection 3.2 uses the particulars of our leading example to illustrate what investigation committed instead simultaneously over both income and environmental concerns could attain.

3.2. Leading Example: What Could Complete Commitment Attain?

As noted in Observation 1, the set of (ex ante) expected utilities that the Sender can obtain under aspect-restricted investigation is a subset of the set achievable under completely committed investigation. Using the details of our main example, for each type $\tau \in \{\tau_1, \tau_2\}$, let me use the short-hand notation $\rho_h^{\tau} \equiv \rho^{\tau}(\overline{a} \mid \theta_h) \in [0,1]$ to identify the probability according to which the decision rule selected by the CEO recommends to launch the new model given that the state of the world is θ_h, for $h = 1, \dots, 9$. Since we already have learned that the CEO has incentives to fully reveal his private information about the pattern of dependence when he is constrained to select a single aspect for investigation, we can aptly consider for comparisons that the Board has also full information about τ when the CEO commits on designing information over both dimensions of uncertainty simultaneously.

Then, when there is no correlation between income growth and environmental concerns ($\tau = \tau_1$), the problem that the CEO must solve to decide about investigation simultaneously over both aspects is:

$$\max_{\{\rho_1^{\tau_1},\dots,\rho_9^{\tau_1}\}} \frac{3}{18}\left[\rho_1^{\tau_1} + \rho_2^{\tau_1} + \rho_3^{\tau_1}\right] + \frac{1}{18}\left[\rho_4^{\tau_1} + \rho_5^{\tau_1} + \rho_6^{\tau_1}\right] + \frac{2}{18}\left[\rho_7^{\tau_1} + \rho_8^{\tau_1} + \rho_9^{\tau_1}\right]$$

$$\text{s.t.:} \qquad 3\rho_3^{\tau_1} + 3\rho_6^{\tau_1} + 2\rho_8^{\tau_1} + 2\rho_9^{\tau_1} \geq 9\rho_1^{\tau_1} + 9\rho_2^{\tau_1} + 3\rho_4^{\tau_1} + 3\rho_5^{\tau_1} + 6\rho_7^{\tau_1}.$$

This is the decision problem that a Sender would need to solve by applying the traditional information design approach to our leading example. In particular, the incentive-compatibility restriction above is simply the expression of the obedience condition (for the presented example) that Bergemann and Morris (2013, 2016, 2019) propose to explore information design. In this case, the CEO would optimally choose $\rho_3^{\tau_1} = \rho_6^{\tau_1} = 2\rho_8^{\tau_1} = \rho_9^{\tau_1} = 1$, any $\rho_1^{\tau_1}, \rho_2^{\tau_1} \in [0,1]$ such that $\rho_1^{\tau_1} + \rho_2^{\tau_1} = 8/9$, and $\rho_4^{\tau_1} = \rho_5^{\tau_1} = 2\rho_7^{\tau_1} = 0$. The CEO would accordingly receive an expected utility of 16/27.

On the other hand, when there is some correlation between income growth and environmental concerns ($\tau = \tau_2$), the problem that the CEO must solve to decide about investigation simultaneously over both aspects is:

$$\max_{\{\rho_1^{\tau_2},\ldots,\rho_9^{\tau_2}\}} \frac{6}{18}\rho_1^{\tau_2} + \frac{2}{18}\rho_2^{\tau_2} + \frac{3}{18}\rho_5^{\tau_2} + \frac{6}{18}\rho_9^{\tau_2}$$
$$\text{s.t.:} \qquad 2\rho_9^{\tau_2} \geq 6\rho_1^{\tau_2} + 3\rho_2^{\tau_2} + 3\rho_5^{\tau_2}.$$

In this case, the CEO would optimally choose $\rho_9^{\tau_2} = 1$, and either (i) any $\rho_2^{\tau_2}, \rho_5^{\tau_2} \in [0,1]$ such that $\rho_2^{\tau_2} + \rho_5^{\tau_2} = 2/3$ and then $\rho_1^{\tau_2} = 0$, or (ii) $\rho_1^{\tau_2} = 1/3$ and then $\rho_2^{\tau_2} = \rho_5^{\tau_2} = 0$. The CEO would accordingly receive an expected utility of 4/9.

Unsurprisingly, the CEO would receive exactly the same expected utility under complete commitment than under aspect-restricted commitment when $\tau = \tau_2$. From the set of states displayed in Fig. 2 (for $\tau = \tau_2$), we can appreciate that the incentives of the CEO and the Board are perfectly aligned under investigation over environmental concerns (the y-aspect), regardless of what investigation about income growth (the x-aspect) can disclose. It follows that investigation only over environmental concerns allows the CEO to place an incentive-compatibility condition that enables him to obtain the same expected utility as under investigation over both aspects combined. However, we also observe that, when $\tau = \tau_1$, investigation over income growth alone delivers an expected utility (16/30) lower than what simultaneous investigation over both aspects allows (16/27), as could be expected in general for a relatively more constrained information designer.

4. Main Results

4.1. Preliminaries

It will be helpful to introduce a couple of ingredients that capture key features of the players' preferences and of the acceptance set of states. Suppose that the true pattern of dependence is some given $\tau \in \mathcal{T}$. For a given aspect $\kappa \in \{x, y\}$, consider then a fixed realization $\kappa_l \in \mathcal{K}$ of the aspect. Let me specify the coefficients:

$$\alpha_\kappa^\tau(\kappa_l) \equiv \sum_{-\kappa \in -\mathcal{K}} \psi_\tau((\kappa_l, -\kappa))\ u_S(\overline{a}, (\kappa_l, -\kappa)) > 0 \qquad \text{and}$$
$$\eta_\kappa^\tau(\kappa_l) \equiv \sum_{-\kappa \in -\mathcal{K}} \psi_\tau((\kappa_l, -\kappa))\ u_R(\overline{a}, (\kappa_l, -\kappa)). \tag{13}$$

The signs of the coefficients $\alpha_\kappa^\tau(\kappa_l)$ and $\eta_\kappa^\tau(\kappa_l)$ capture whether or not, respectively, the Sender and the Receiver would prefer the acceptance action $\overline{a}$ if they considered the given priors

ψ_τ and knew that the realization of aspect κ is κ_l. Furthermore, the size of such coefficients is informative of how much the players like (or dislike) the acceptance action under such conditions. The model places little structure on the signs of the coefficients $\{\eta^\tau_\kappa(\kappa_l)\}_{l=1}^m$. In general, the signs of such coefficients depend on the shape of the acceptance set $\overline{\Theta}$ and on the corresponding prior ψ_τ about the state. From Assumption 2 (iv) on the players' preferences, we observe nonetheless that $\sum_{l=1}^m \eta^\tau_\kappa(\kappa_l) = \sum_{\theta\in\Theta} \psi_\tau(\theta)\ u_R(\overline{a}_S, \theta) < 0$. Therefore, we know that the condition $\min_{\kappa_l\in\mathcal{K}}\ \eta^\tau_\kappa(\kappa_l) < 0$ must be satisfied. In other words, since the preferences of the Sender are such that $\alpha^\tau_\kappa(\kappa_l) > 0$ for each $\kappa_l \in \mathcal{K}$, the model assumes that, for each possible pattern of dependence τ, there is always at least one realization of each aspect choice κ, which belongs to the set $\arg min_{\kappa_l\in\mathcal{K}}\ \eta^\tau_\kappa(\kappa_l)$, such that Sender and Receiver disagree completely over the suitability of accepting the proposal. Theorem 1 then shows that, for the Sender to be able to persuade the Receiver through information design, there must exist some other realization, for at least one of the two aspects, over which the two players fully agree that the Sender's ideal action $\overline{a}$ is indeed the best course of action.

Turning to the Sender's ex ante optimization problem specified in Eq. (9), consider the case where the Sender decides to fully reveal his private information about the pattern of dependence ($d = r$). Notice that by applying Assumption 2 (iv) regarding the Sender's preferences, it must be the case that we have $\sigma^{(\tau,r)^*}_\kappa(\underline{a} \mid \kappa_l) = 0$ in equilibrium, for each given realization $\kappa_l \in \mathcal{K}$. Given an aspect choice $\kappa \in \mathcal{K}$ and a type $\tau \in \mathcal{T}$, the problem of the Sender expressed in Eq. (9), for the fully revealing message strategy $d = r$, can then be rewritten as the linear programming problem:

$$[\mathcal{P}] \qquad \begin{aligned} \max_{\{\hat{\sigma}^\tau_\kappa\in[0,1]^m\}} \quad & \sum_{l=1}^m \alpha^\tau_\kappa(\kappa_l)\ \hat{\sigma}^\tau_{\kappa_l} \\ \text{s.t.:} \quad & \sum_{l=1}^m \eta^\tau_\kappa(\kappa_l)\ \hat{\sigma}^\tau_{\kappa_l} \geq 0. \end{aligned} \tag{14}$$

Under the description of the Sender's problem $[\mathcal{P}]$ provided by Eq. (14) above, the fact that $\alpha^\tau_\kappa(\kappa_l) > 0$ for each aspect choice κ and each possible pattern of dependence τ guarantees a well-defined objective function $\sum_{l=1}^m \alpha^\tau_\kappa(\kappa_l)\ \hat{\sigma}^\tau_{\kappa_l}$ for problem $[\mathcal{P}]$ to have a solution. Yet, to ensure that problem $[\mathcal{P}]$ has in fact a solution where the Sender is able to influence the Receiver to accept, we need to verify that the incentive compatibility condition $\sum_{l=1}^m \eta^\tau_\kappa(\kappa_l)\ \hat{\sigma}^\tau_{\kappa_l} \geq 0$ of problem $[\mathcal{P}]$ allows indeed for solutions where the Sender's ideal action $\overline{a}$ is optimally recommended with positive probability.

4.2. Information Design under Full Revelation of the Pattern of Dependence

Theorem 1 gives us a set of conditions on the possible priors about the state and on the players' preferences that characterize the optimal information design behavior that the Sender follows to persuade the Receiver, provided that he fully reveals his private information about the pattern of dependence. The proofs of all the results of the paper are relegated to the Appendix.

Theorem 1. *Let Assumption 1 and Assumption 2 hold. Suppose that the Sender selects aspect $\kappa \in \mathcal{K}$ to make investigation choices over it and chooses to fully reveal his private information about the pattern of dependence ($d = r$). Then, (a) for a given type $\tau \in \mathcal{T}$, investigation over aspect κ is able to persuade the Receiver to accept the proposal if and only if the set of aspect realizations $\hat{\mathcal{K}}^\tau \equiv \{ \hat{\kappa}^\tau \in \mathcal{K} \mid \eta^\tau_\kappa(\hat{\kappa}^\tau) > 0 \}$ is nonempty. Provided that the set of aspect realizations $\hat{\mathcal{K}}^\tau$ is nonempty, then (b) investigation optimally recommends acceptance of the proposal with probability one for each realization $\hat{\kappa}^\tau \in \hat{\mathcal{K}}^\tau$ whereas, for the remaining realizations $\kappa^\tau \in \mathcal{K} \setminus \hat{\mathcal{K}}^\tau$ of the aspect, investigation recommends acceptance with positive probability (less than one) only for those realizations $\bar{\kappa}^\tau \in \mathcal{A}^\tau_\kappa \subseteq \mathcal{K} \setminus \hat{\mathcal{K}}^\tau$, where*

$$\mathcal{A}^\tau_\kappa \equiv arg \max_{\kappa^\tau \in \mathcal{K} \setminus \hat{\mathcal{K}}^\tau} \quad \frac{\alpha^\tau_\kappa(\kappa^\tau)}{\eta^\tau_\kappa(\kappa^\tau)}.$$

In the case where both sets $\hat{X}^\tau$ and $\hat{Y}^\tau$ are nonempty—so that the Sender can persuade the Receiver by choosing investigation processes over any of the two aspects of uncertainty—, (c) the Sender picks arbitrarily one aspect realization $\bar{\kappa}^\tau \in \mathcal{A}^\tau_\kappa$ for each of the two aspects $\kappa \in \{x, y\}$ and then optimally select(s) for investigation the aspect(s) κ^* that solve(s) the problem

$$\max_{\kappa \in \{x,y\}} \quad \sum_{\hat{\kappa}^\tau \in \hat{\mathcal{K}}^\tau} \Big[\frac{\alpha^\tau_\kappa(\bar{\kappa}^\tau)}{\eta^\tau_\kappa(\bar{\kappa}^\tau)} \eta^\tau_\kappa(\hat{\kappa}^\tau) + \alpha^\tau_\kappa(\hat{\kappa}^\tau) \Big].$$

The condition (a) identified by Theorem 1 ensures that the Sender's problem $[\mathcal{P}]$ has solutions where the chosen information structures recommend acceptance with probability one, conditional on some aspect realization(s). In particular, the condition that the set $\hat{\mathcal{K}}^\tau$ be nonempty means that there is always at least one realization $\hat{\kappa}^\tau$ of the chosen aspect κ over which the interests of the two players are perfectly aligned for the Sender's ideal action to be recommended following the required incentive-compatibility condition. We can then naturally refer to the set $\hat{\mathcal{K}}^\tau$ as the *agreement set of realizations for aspect* κ, provided that the true pattern of dependence is τ. Then, if the Sender optimally recommends acceptance with probability one for some realizations of the chosen aspect, result (b) of Theorem 1 gives us the criterion under which the Sender recommends acceptance with non zero (yet, less than one) probability, conditional on aspect realizations other than the ones in the agreement set. Lastly, for those cases where the primitives of the model enable the Sender to persuade the Receiver to accept by designing information over any of the two aspects of uncertainty, condition (c) of Theorem 1 describes the criterion that determines his optimal aspect choice for information design.

A case of particular interest in many applications is the one that arises when the preferences of the Sender are constant across all possible states, and the preferences of the Receiver do not change across states within the acceptance set, on the one side, and across all states outside the acceptance set, on the other side. The following Assumption 4 captures these situations. In particular, the preferences of the players in our leading example of Subsection 2.6 satisfy Assumption 4.

Assumption 4. *The preferences of the players satisfy:*

(i) $u_S(\underline{a}, \theta) = \overline{u}_S > 0$ *for each* $\theta \in \Theta$*;*

(ii) $u_R(\overline{a}, \theta) = \overline{u}_R > 0$ *if* $\theta \in \overline{\Theta}$ *and* $u_R(\overline{a}, \theta) = \underline{u}_R < 0$ *if* $\theta \in \Theta \setminus \overline{\Theta}$.

For an aspect $\kappa \in \{x, y\}$ and a given aspect realization $\bar{\kappa} \in \mathcal{K}$, we will be interested in considering a subset of aspect realizations $-\overline{\mathcal{K}}(\bar{\kappa}) \subseteq -\mathcal{K}$ for the remaining aspect $-\kappa$, which crucially depends on the shape of the acceptance set $\overline{\Theta}$ according to the specification:

$$-\overline{\mathcal{K}}(\bar{\kappa}) \equiv \{-\kappa \in -\mathcal{K} \mid (\bar{\kappa}, -\kappa) \in \overline{\Theta}\}.$$

Observe that $\overline{\Theta} = \cup_{\bar{\kappa} \in \mathcal{K}} - \overline{\mathcal{K}}(\bar{\kappa})$. The following Corollary 1 to Theorem 1 describes how the Sender designs information under the class of preferences given by Assumption 4, provided that he fully reveals his private information about the pattern of dependence. Corollary 1 follows directly by applying the results of Theorem 1 to an environment under the preference specification in Assumption 4, upon use of the coefficients specified earlier in Eq. (13).

Corollary 1. *Let Assumption 1, Assumption 2, and Assumption 4 hold. Suppose that the Sender selects aspect* $\kappa \in \mathcal{K}$ *to make investigation choices over it and chooses to fully reveal his private information about the pattern of dependence* ($d = r$)*. Then, (a) for a given type*

$\tau \in \mathcal{T}$, investigation over aspect κ is able to persuade the Receiver to accept the proposal if and only if the set

$$\hat{\mathcal{K}}^{\tau} = \left\{ \hat{\kappa}^{\tau} \in \mathcal{K} \;\middle|\; \frac{\sum_{-\kappa \in -\overline{\mathcal{K}}(\hat{\kappa}^{\tau})} \psi_{\tau}((\hat{\kappa}^{\tau}, -\kappa))}{\psi_{\kappa}(\hat{\kappa}^{\tau})} > \frac{-\underline{u}_R}{\overline{u}_R - \underline{u}_R} \right\}$$

is nonempty. Provided that the set of aspect realizations $\hat{\mathcal{K}}^{\tau}$ is nonempty, then (b) investigation optimally recommends acceptance of the proposal with probability one for each realization $\hat{\kappa}^{\tau} \in \hat{\mathcal{K}}^{\tau}$ whereas, for the remaining realizations of the aspect $\kappa^{\tau} \in \mathcal{K} \setminus \hat{\mathcal{K}}^{\tau}$, investigation recommends acceptance with positive probability (less than one) only for those realizations $\bar{\kappa}^{\tau} \in \mathcal{A}^{\tau}_{\kappa} \subseteq \mathcal{K} \setminus \hat{\mathcal{K}}^{\tau}$, such that

$$\mathcal{A}^{\tau}_{\kappa} \equiv arg \max_{\kappa^{\tau} \in \mathcal{K} \setminus \hat{\mathcal{K}}^{\tau}} \quad \overline{u}_S \Bigg/ \left[\frac{\underline{u}_R}{\overline{u}_R - \underline{u}_R} + \mathbb{I}_{\overline{\Theta}}(\kappa^{\tau}) \frac{\sum_{-\kappa \in -\overline{\mathcal{K}}(\hat{\kappa}^{\tau})} \psi_{\tau}((\kappa^{\tau}, -\kappa))}{\psi_{\kappa}(\kappa^{\tau})} \right],$$

where $\mathbb{I}_{\overline{\Theta}}$ is the indicator function defined as $\mathbb{I}_{\overline{\Theta}}(\kappa^{\tau}) = 1$ if $\kappa^{\tau} \in \overline{\mathcal{K}}$ and $\mathbb{I}_{\overline{\Theta}}(\kappa^{\tau}) = 0$ if $\kappa^{\tau} \in \mathcal{K} \setminus \overline{\mathcal{K}}$. In the case where both sets $\hat{X}^{\tau}$ and $\hat{Y}^{\tau}$ are nonempty—so that the Sender can persuade the Receiver by choosing investigation processes over any of the two aspects of uncertainty—, (c) the Sender picks arbitrarily one aspect realization $\bar{\kappa}^{\tau} \in \mathcal{A}^{\tau}_{\kappa}$ for each of the two aspects $\kappa \in \{x, y\}$ and then optimally selects the aspect(s) κ^* that solve(s) the problem

$$\max_{\kappa \in \{x,y\}} \sum_{\hat{\kappa}^{\tau} \in \hat{\mathcal{K}}^{\tau}} \chi_{\psi_{\tau}}(\hat{\kappa}^{\tau}; \bar{\kappa}^{\tau}),$$

where, for each given pattern of dependence $\tau \in \mathcal{T}$, and each given pair of aspect realizations $\hat{\kappa}^{\tau} \in \hat{\mathcal{K}}_{\tau}$ and $\bar{\kappa}^{\tau} \in A^{\tau}_{\kappa}$, the term $\chi_{\tau}(\hat{\kappa}^{\tau}; \bar{\kappa}^{\tau})$ is specified as:

$$\chi_{\tau}(\hat{\kappa}^{\tau}; \bar{\kappa}^{\tau}) \equiv \psi_{\kappa}(\hat{\kappa}^{\tau}) + \frac{\underline{u}_R \psi_{\kappa}(\hat{\kappa}^{\tau}) + (\overline{u}_R - \underline{u}_R) \sum_{-\kappa \in -\overline{\mathcal{K}}(\hat{\kappa}^{\tau})} \psi_{\tau}((\hat{\kappa}^{\tau}, -\kappa))}{\underline{u}_R + (\overline{u}_R - \underline{u}_R) \sum_{-\kappa \in -\overline{\mathcal{K}}(\bar{\kappa}^{\tau})} \psi_{\tau}((\bar{\kappa}^{\tau}, -\kappa))/\psi_{\kappa}(\bar{\kappa}^{\tau})}.$$

In short, the insights of Corollary 1 allow us to determine how the Sender would optimally choose investigation, for each possible pattern of dependence, basically by studying the shape of the priors about the state for the corresponding pattern. Although the expressions derived are lengthy and seem complex, the insights of Corollary 1 allows us to use quite readily the characterization result provided by Theorem 1 in applications. Let me illustrate this point by applying the result of Corollary 1 to our leading example.

4.3. Leading Example: Simplified Analysis (Using Corollary 1)

By resorting to Corollary 1, let me apply the characterization provided by Theorem 1 to our main example of Subsection 2.6. First, recall that $\overline{u}_S = 1$, $\underline{u}_R = -1$, and $\overline{u}_R = 1/3$. Therefore, $-\underline{u}_R/(\overline{u}_R - \underline{u}_R) = 3/4$. Also, notice that $\overline{X}(y_1) = \emptyset$, $\overline{X}(y_2) = \{x_3\}$, and $\overline{X}(y_3) = \{x_1, x_2, x_3\}$, whereas $\overline{Y}(x_1) = \{y_3\}$, $\overline{Y}(x_2) = \{y_3\}$, and $\overline{Y}(x_3) = \{y_2, y_3\}$.

Suppose first that the CEO's true type is $\tau = \tau_1$ so that income growth and environmental concerns are unrelated with each other. Then, for investigation over aspect $\kappa = x$, we can readily compute:

$$\frac{\sum_{y\in\overline{Y}(x_1)} \psi_{\tau_1}((x_1, y))}{\psi_x(x_1)} = 1/3, \ \frac{\sum_{y\in\overline{Y}(x_2)} \psi_{\tau_1}((x_2, y))}{\psi_x(x_2)} = 1/3, \ \text{and} \ \frac{\sum_{y\in\overline{Y}(x_3)} \psi_{\tau_1}((x_3, y))}{\psi_x(x_3)} = 2/3,$$

so that $\hat{X}^{\tau_1} = \emptyset$. On the other hand, for investigation over aspect $\kappa = y$, we obtain

$$\frac{\sum_{x\in\overline{X}(y_1)} \psi_{\tau_1}((x, y_1))}{\psi_y(y_1)} = 0, \ \frac{\sum_{x\in\overline{X}(y_2)} \psi_{\tau_1}((x, y_2))}{\psi_y(y_2)} = 1/3, \ \text{and} \ \frac{\sum_{x\in\overline{X}(y_3)} \psi_{\tau_1}((x, y_3))}{\psi_y(y_3)} = 1,$$

so that $\hat{Y}^{\tau_1} = \{y_3\}$. Using the insights from Corollary 1, we observe that investigation over aspect x is useless to persuade the Board, whereas investigation over aspect y is capable of persuading her by recommending acceptance with probability one when the realization of aspect y is y_3. In addition, for the case of investigation over aspect y, we easily observe that $\mathcal{A}_y^{\tau_1} = \{y_2\}$. Thus, the CEO's investigation decision can optimally recommend acceptance with positive probability for realization y_2 as well.

Now, suppose that the CEO's true type is $\tau = \tau_2$ so that there is some correlation between income growth and environmental concerns. Then, for investigation over aspect $\kappa = x$, we can compute:

$$\frac{\sum_{y\in\overline{Y}(x_1)} \psi_{\tau_2}((x_1, y))}{\psi_x(x_1)} = 0, \ \frac{\sum_{y\in\overline{Y}(x_2)} \psi_{\tau_2}((x_2, y))}{\psi_x(x_2)} = 0, \ \text{and} \ \frac{\sum_{y\in\overline{Y}(x_3)} \psi_{\tau_2}((x_3, y))}{\psi_x(x_3)} = 1,$$

so that $\hat{X}^{\tau_1} = \{x_3\}$. On the other hand, for investigation over aspect $\kappa = y$, we obtain

$$\frac{\sum_{x\in\overline{X}(y_1)} \psi_{\tau_2}((x, y_1))}{\psi_y(y_1)} = 0, \ \frac{\sum_{x\in\overline{X}(y_2)} \psi_{\tau_2}((x, y_2))}{\psi_y(y_2)} = 0, \ \text{and} \ \frac{\sum_{x\in\overline{X}(y_3)} \psi_{\tau_2}((x, y_3))}{\psi_y(y_3)} = 1,$$

so that $\hat{Y}^{\tau_1} = \{y_3\}$. Using the insights from Corollary 1, we observe that investigation over aspect x is capable of persuading the Board upon recommending acceptance with probability one when the realization of aspect x is x_3. Also, investigation over aspect y is capable of persuading the Board by recommending acceptance with probability one when the realization of aspect y is y_3. For the case of investigation over aspect x, we trivially obtain that $\mathcal{A}_x^{\tau_2} = \{x_1, x_2\}$. Therefore, acceptance can optimally be recommended with positive probability for any of the realizations x_1 and/or x_2, with the restriction imposed the incentive-compatibility constraint (with equality). Similarly, for the case of information design over aspect y, we easily observe that $\mathcal{A}_y^{\tau_2} = \{y_1, y_2\}$.

Finally, to determine which aspect the CEO optimally selects for information design, for each pattern of dependence, notice that

- For type $\tau = \tau_1$, the CEO optimally chooses aspect $k = y$, since $\hat{X}^{\tau_1} = \emptyset$ whereas $\hat{Y}^{\tau_1} = \{y_3\}$.
- For type $\tau = \tau_2$, we obtain

$$\chi_{\tau_2}(x_3; x_1) = \chi_{\tau_2}(x_3; x_2) = \frac{(-1)(2/6) + (1/3 + 1)(2/6)}{(-1) + (1/3 + 1)(0)} + (1)(2/6) = 2/9,$$

whereas

$$\chi_{\tau_2}(y_3; y_1) = \chi_{\tau_2}(y_3; y_2) = \frac{(-1)(1/3) + (1/3 + 1)(2/6)}{(-1) + (1/3 + 1)(0)} + (1)(1/3) = 2/9,$$

so that the CEO of type τ_2 is optimally indifferent between choosing any aspect, either x or y, for investigation.

We observe that the entire description provided now for the CEO's optimal disclosure behavior coincides exactly with the one presented earlier in Subsection 2.6.

4.4. Full Revelation about the Pattern of Dependence

Theorem 1 has characterized how the Sender optimally designs information in order to persuade the Receiver towards acceptance, yet conditional on the Sender optimally revealing his private information about the pattern of dependence. Therefore, we need to address the question of whether the Sender has in fact incentives to reveal his private information about τ in a way consistent with the induced posterior beliefs of the Receiver, given her skeptical revision process. The insight that emerged from the leading example that in equilibrium the

Sender fully reveals his private information about the relationships between the two aspects, upon designing information over one aspect, is a fairly general implication in the proposed framework. Theorem 2 provides the formal statement of this result, under the condition that the Receiver is maximally skeptical according to the criterion in Assumption 3.

Theorem 2. *Let the proposed information disclosure game satisfy Assumption 1, Assumption 2, andAssumption 3. Then, in each equilibrium* $\phi^* = ((\sigma_x^*, \sigma_y^*); d^*, \kappa^*)$ *and* β^{d^*} *of the proposed disclosure game, the Sender fully reveals his private information about the relationships between the two aspects (i.e.,* $d^* = r$*).*

In the proposed benchmark skepticism plays a role essentially different to the one that underlies the classical "unravelling" mechanism of the persuasion literature. Suppose that the Sender chooses to withhold some of his private information about the pattern of dependence. Then, optimal aspect-restricted information design applied to environments with two available actions (either accept or reject), leaves the skeptical Receiver indifferent between the various types that are pooled under a common message. Intuitively, information design disciplines the Receiver so that the original conflict of interests is crucially lessen. This is not the case in the traditional verifiable disclosure model, wherein both Sender and Receiver rank the set of possible types in different directions and there is always a sharp conflict of interests described by such opposed rankings. On the other hand, the skeptical Receiver in the current model places positive probability, according to her priors, on each of the Sender's possible types that are pooled together. In addition to this, it follows from the (linear) structure of the information design problem that the (ex ante) expected utility that accrues to the Sender upon pooling about his type can be expressed as a convex combination of the expected utilities that he would receive upon full revelation. Conditional on the Receiver accepting the proposal $\overline{a}$, such an implication gives strict incentives to the type associated to the Sender's highest expected utility to separate from the rest of types that are grouped together under the common message. Then, through a recursive process, the type that is left in each iteration to provide the Sender with the highest expected utility always wishes to separate itself from the types that remain pooled under a common message. By iterating, this mechanism leads to the full revelation result stated in Theorem 2.

5. Application: Media *Slant*

I am particularly motivated by applications that allow us to understand better media tactics of persuasion. Under the restriction that investigation over all aspects of uncertainty is not possible before the relevant decision must be made, the proposed framework offers a logic for how media outlets *slant*[29] by strategically selecting ex ante investigation and interim communication. One strand of the economic literature on media slant considers that either Receivers (Mullainathan and Shleifer, 2005), Senders (Baron, 2006), or public institutions (Besley and Prat, 2006) are biased in ways such that they benefit from the Senders distorting the information they provide. Other contemporary efforts to explain media slant consider a reputation motive that gives Senders the incentives to distort their reporting (Gentzkow and Shapiro, 2006). Unlike these papers, neither any sort of exogenous preference for distorted news nor reputation concerns are required to obtain a rationale for slant under the setup proposed in this paper.

I will use an example proposed by Mullainathan and Shleifer (2005) to illustrate how the current model could be useful to interpret and rationalize media slant. Suppose that the Bureau of Labor Statistics (BLS) discloses that the number of unemployed increased by 200,000, from 6.1 percent to a 6.3 unemployment rate. We can interpret this finding as the outcome of committed investigation about a particular economic aspect of interest to decision makers (Receivers). Mullainathan and Shleifer (2005) stress that, alongside with the outcome of such an investigation, media outlets can offer two alternative headlines that "interpret" further the information provided by the BLS. Together with the new data, one headline, "Recession Fears Grow," can suggest an imminent recession. This headline also offers the views of Harvard economist John K. Galbraith, who sees this as a sign of inadequate economic policies. In particular, he says "not since Herbert Hoover has a president ignored economic realities so blatantly." Another headline, "Turnaround in Sight," can point instead towards an imminent expansion. This second headline offers the views of the chief stock market analyst of Goldman and Sachs, Abbie J. Cohen, who sees this as a sign of profitable investment opportunities. She says that "this is a good time to increase exposure to stocks both because of the strong underlying fundamentals and the softness in the labor market bodes well for corporate profitability."

Using the benchmark proposed here, one could formalize this situation by considering three relevant aspects of uncertainty for the Receiver.[30] A state of the world in this case could be $\theta = (x, y, z)$, where x describes job market conditions, y accounts for the quality

[29] The idea of *slant* was introduced by Hayakawa (1940). Slant roughly refers to the act of selecting verifiable details that are favorable or unfavorable to a certain view for an action choice under uncertainty.

[30] Although the model has been developed in terms of two dimensions of uncertainty for simplicity, its functioning and implications go through for a general multi-dimensional state with a finite number of coordinates.

of economic policies, and z measures the position of economic fundamentals relevant for the stock market. Suppose that, as an expert on the quality of economic policies, J. K. Galbraith owns private information about whether or not the implemented policies are adequate. Similarly, suppose that, as an expert on economic fundamentals for corporate performance, A. J. Cohen holds private information about the profitability of investments in the stock market. Such pieces of private information possessed by each of the two experts can even have the form of verifiable reports (perhaps obtained from previous research). We can consider that the Receiver in this application is an undecided voter that must choose whether or not to support the current governing party for the next term. Suppose that the media outlet wishes the voter to support the current governing party, regardless of the state of the world. We can consider that this voter is able to pay attention to data from investigation about any single one of the relevant aspects. Similarly, we can also consider that the media outlet can commission investigation about any single one of the aspects before election day.

If we wish to interpret this story according to the insights of the proposed model, then the preferences of the Receiver are crucial to determine how a media outlet will combine the data released by the BLS with its (interim) strategic communication. One possibility is that the voter relies only on her beliefs about job market conditions and about the quality of the economic policies implemented by the current government (say, Receiver A). Another possibility, though, is that the voter instead casts her vote based only on her beliefs about job market conditions and on the investment opportunities available in the economy (say, Receiver B). We can capture these considerations by considering two different acceptance sets, $\overline{\Theta}_A$ and $\overline{\Theta}_B$, such that the respective Receiver either does not care whatsoever about the z-aspect ($\overline{\Theta}_A$), or does not care at all about the y-aspect ($\overline{\Theta}_B$). Then, the crucial point highlighted in the proposed model is that the media outlet chooses the aspect over which to commit to investigation based on the preferences of the Receiver. In particular, the proposed model would predict that the media indeed commits to disclose data about unemployment (obtained by the BLS's investigation) if and only if there is at least one possible unemployment rate conditional on which the media firm and the respective Receiver totally agree on that supporting the governing party is the right action. In addition, if the media chooses to commit on investigation only about job market conditions, then the result of Theorem 1 would translate into that, for any of the two possible Receivers, A or B, the incentives of the media and of the respective Receiver are relatively more aligned when evidence is disclosed about job conditions. Intuitively, taking the average over all possible observations for any of the three relevant aspects, both Receivers are more willing to vote for the governing party based on observations of employment data, relative to data about economic policy quality or about fundamentals of the stock market. If these conditions are satisfied, the implications of Theorem 2 would translate into that the media outlet chooses to reveal all the available private information, either about the relationships between job market conditions and quality of economic policies (when the Receiver is A) or about relationships between job market conditions and investment opportunities in the stock market (when the Receiver is B). Offering the views of renowned experts such as either J. K. Galbraith or A. J. Cohen, depending on each of the respective patterns of dependence, is then a natural approach in this story that seems consistent with the implications of the proposed model.

The logic that the proposed model provides for rationalizing media slant is consistent with other theoretical proposals and with empirical findings as well. Although Mullainathan and Shleifer (2005)'s behavioral assumption that Receivers want to see their initial beliefs confirmed makes the analyses quite different, there are also similarities in the ways in which the implications about slant work. In particular, they find that Senders slant by disclosing information in order to adjust to the tastes of the Receivers, which shares the basic rationale of the result that the Sender chooses a single aspect for investigation in a way crucially driven the preferences of the Receiver. Interestingly, at the empirical level, Gentzkow and Shapiro (2010) estimate that roughly a 20 of the variation on their recent US sample on media slant obeys to media outlets having incentives to respond to the preferences of their consumers. In their estimations, the identities, or preferences, of the media outlets play no role in explaining slant. In my model, the Sender "tailors" his aspect-restricted information design choices to the Receiver's tastes while, at the same time, informs her completely about what he knows of the relationships between the relevant aspects.

6. Further Literature Connections

The study of influential communication goes back to the literature on strategic revelation, or advise, from informed Senders to uninformed Receivers. Starting with the seminal contributions of Green and Stokey (1980) and Crawford and Sobel (1982), the cheap talk framework establishes that influential communication is critically bounded when the conflict of interests is high.[31] Credible communication is enhanced when the Sender can commit to design

[31] Although fully revealing cheap talk communication can be obtained—under certain conditions—when there are multiple dimensions of uncertainty (Battaglini, 2002; Chakraborty and Harbaugh, 2007, 2010), the message that costless−uncommitted−unverifiable disclosure is severely restricted continues to hold even for such environments. In particular, multi-dimensional cheap talk communication is severely restricted if either

information.[32] Following the information design (Bergemann and Morris, 2013, 2016)—or Bayesian persuasion (Rayo and Segal, 2010; Kamenica and Gentzkow, 2011)—approach, the current paper has considered that the Sender has the ability to make (ex ante) commitments to design information. In contrast to this literature, however, the setup proposed here assumes that the Sender lacks the ability to commit over all dimensions of uncertainty simultaneously, neither can he commit over possible correlations, and must instead resort to full, but "isolated," commitment over any single one of the separate dimensions.

The framework here proposed builds also upon the influential contributions on verifiable information disclosure (Grossman, 1981; Milgrom, 1981; Milgrom and Roberts, 1986; Krishna and Morgan, 2001; Che and Kartik, 2009; Kartik et al., 2017). Within the verifiable disclosure literature, there are clear connections, mainly in motivation and raised questions, to models that deal with how selective reporting in multi-dimensional environments can affect security price dynamics (Shin, 1994, 2003). Gentzkow and Kamenica (2017c) enrich the classical verifiable disclosure setup by adding a stage where the Sender, prior to choosing how much verifiable information he discloses, commits to costly obtaining the relevant information. They show that there is always an equilibrium where all the acquired information is disclosed and that, if there is another equilibrium where information is withheld, then the Sender must receive the same expected utility than in the full disclosure equilibrium. This result bears a clear resemblance with the full disclosure result that the setup proposed here delivers. Their consideration that the Sender combines committed information acquisition with verifiable disclosure is also shared in the current setup. The models are substantially different though as my interest is in multi-dimensional environments where the Sender fully commits to disclose information about one of the aspects and, in addition, can strategically choose either unverifiable or verifiable disclosure about how the aspects correlate.

Since one natural motivation for the restriction that the Sender must choose information structures over separate dimensions of uncertainty is to consider that the Receiver is able to process information from investigation about any single one of the aspects, this paper is also related to the study of persuasion rules under restrictions on the amount of verifiable information that can be accumulated or processed that was explored by Glazer and Rubinstein (2004).

---

the state space is bounded (Ambrus and Takahashi, 2008) or if the dimensions of uncertainty are strongly correlated between them according to the prior (Levy and Razin, 2007).

[32] As already suggested by Levy and Razin (2007)'s analysis even before the Bayesian persuasion literature comprehensively explored these topics.

The pursued approach of weakening the Sender's commitment power bears also similarities with the setup proposed by Nguyen and Tan (2019). They consider a Bayesian persuasion model where the Sender first receives privately the signals disclosed by the selected information structure and then sends a costly message to the Receiver about the original signals. In addition to the analytical similarities, there is also a resemblance in the stories that the settings try to rationalize: in both papers, the Sender somehow strategically "interprets" the information disclosed by a committed information structure. The differences, though, are also of substance. Nguyen and Tan (2019) do not consider information design restricted to particular aspects of uncertainty and they model the communication that complements the information structure as costly signaling. In the current paper, such an additional source of information has the form of costless communication.

Sharing motivations and interests with Nguyen and Tan (2019), other papers have also weaken the commitment power of the Sender in information design problems. In Min (2017) and Frechette et al. (2019), the Sender has the ability to distort the information structure that he selects before releasing it to the Receiver. Lipnowski and Ravid (2019) relaxes the Sender's commitment power to formalize cheap talk environments where the payoffs of the Sender do not depend on the state of the world, while Lipnowski et al. (2019) analyses credibility issues regarding the quality of the chosen information structures. At a more generally level, Bester and Strausz (2001) have explored how incentive-compatibility conditions and the Revelation Principle need to be modified in contracting situations with a single agent where the principal has limited commitment power. Importantly, all these papers are part of a recent literature on persuasion where the key consideration is *limited or partial commitment*: there is an exogenous positive probability that the chosen information structure is not binding for the information designer. Unlike this assumption, in the current paper the Sender can commit fully, without any possibility of subsequent manipulation of the chosen information structure. The crucial consideration, which I term as *aspect-restricted commitment*, is that he can only commit to one of the dimensions of a two-dimensional relevant uncertainty parameter.

For environments where the relevant uncertainty is multi-dimensional, Frankel and Kartik (2019) consider a model of separate costly signaling over two aspects of uncertainty and obtain that information provision about one aspect can diminish ("muddle") disclosure over the other aspect. Unlike their paper, though, the current setup does not consider costly signaling.

Similar to the assumption that the Sender chooses to design information over one aspect or another of uncertainty, Deimen and Szalay (2019) allow the Sender to acquire information selectively about one aspect (the Sender's ideal action) or another (the Receiver's ideal action). Theirs is a setup of delegated expertise enriched with the possibility costless information acquisition about two separate aspects of uncertainty and cheap talk communication. Other than the shared interest in strategic aspect choice for information provision, the setup and questions explored are quite distinct.

Also, the questions investigated in the current paper are reminiscent of those addressed by Che et al. (2013), who explore how verifiable ("hard") information in the hands of the Receiver combines with comparative cheap talk ("soft information") by the Sender. They obtain that, in equilibrium, the Sender biases his cheap talk towards recommendations favorable to the Receiver ("pandering"), provided that there is a mild conflict of interests only over an outside option. The current paper focuses less on the interplay between verifiable and unverifiable information, and considers instead the interplay between ex ante and interim information decisions.

At a more instrumental level, the maintained assumption that the Receiver begins with a non-fully identified probabilistic model and then uses beliefs over prior distributions is closely related to the *multiple priors* decision benchmark that was axiomatized by Gilboa and Schmeidler (1989) and subsequently investigated by a number of papers concerned about the role of *ambiguity* in preference representation under uncertainty (Bewley, 2002; Klibanoff et al., 2005; Maccheroni et al., 2006a,b; Seo, 2009). More recently, Acemoglu et al. (2016) have also made use of non-fully identified priors to explore the robustness of asymptotic agreement in learning environments.

7. Concluding Comments

This paper has explored persuasion by information disclosure for environments with two-aspect relevant uncertainty. The novelty of the approach lies in the assumption that the Sender is constrained to combine (ex ante) investigation over any single one of the aspects and (interim) strategic communication about the dependencies between the two aspects. The equilibrium notion considers that the Receiver is highly skeptical when revising her priors about the dependencies between the aspects. In particular, upon communication choices that do not allow the Receiver to learn the dependencies, she considers that the true dependencies are the least beneficial ones for her. Under this natural criterion, equilibria feature full revelation of the private information that the Sender has about how the two aspects are related. Given this full revelation result, the Sender commits to disclose data from investigation for the aspect associated to possible observations that, in average, mitigate the most the initial conflict of interests. For persuasive aspect-restricted investigation to take place, it is necessary (and sufficient) to have at least one piece of data about the chosen aspect over which the incentives of the two parties are perfectly aligned. Under the usual maximal skepticism approach to strategic communication, the full revelation insight stems precisely from the interaction of (ex ante) investigation over only one of the aspects and such (interim) communication decisions.

This paper restricted attention to two-aspect settings where commitment for investigation is feasible for only one of the two aspects. While this is convenient for tractability and expositional reasons, the mechanisms that drive the model's qualitative implications go entirely through for settings with general (finite) multi-dimensional uncertainty and commitment restricted to (strict) subsets of all the relevant aspects. For such cases, following the reasoning behind the core result of Theorem 1, it is intuitive to see that the Sender would analogously select the subset of aspects for investigation whose possible realizations alleviate (in average) the most the original source of conflict of interests. In addition, each optimally selected investigation would similarly recommend higher acceptance probabilities the lower the conflict of interests associated with the corresponding aspect realization. The basic logic of the key result of full revelation of the pattern of dependence also extends intuitively to more general settings. In such cases, the Sender would reveal fully his private information about the correlations between an "investigated aspect" and each other "non-investigated aspect." This revelation behavior would be supported by the implication that investigation on a given aspect disciplines the Receiver in a way such that—upon the event of no full revelation of the aspect—she is left indifferent among the various possible relationships between such an aspect and any other "non-investigated aspect."

The full credible communication insight rests crucially on the consideration that, when left indifferent between several pooled types of the Sender, the Receiver resorts, in a proportional way, to her initial beliefs. This guarantees that any possible type is assigned a positive probability of occurrence, which always triggers the incentives of some pooled type to separate from the rest. In spite of not being the unique proposal one can conceivably make for these cases, it seems to describe the most neutral approach that the Receiver can take when left indifferent among several unknown types of the Sender.

8. Appendix

Omitted Proofs

Proof of Theorem 1.
Suppose that the Sender selects an aspect $\kappa \in \{x, y\}$ to commission investigation and chooses to fully reveal his private information about the pattern of dependence $(d = r)$.

Consider first the case where, for each pattern of dependence $\tau \in \mathcal{T}$, the set of aspect realizations $\hat{\mathcal{K}}^\tau = \{\ \hat{\kappa}^\tau \in \mathcal{K} \mid \eta^\tau_\kappa(\hat{\kappa}^\tau) > 0\ \}$ is empty so that $\eta^\tau_\kappa(\kappa_l) \leq 0$ for each aspect realization $\kappa_l \in \mathcal{K}$. Then, we observe that the only decision rule $\hat{\sigma}^\tau_\kappa = \{\hat{\sigma}^\tau_{\kappa_l} \in [0,1]\}_{l=1}^m$ that solves the Sender's problem $[\mathcal{P}]$ when he has type τ (detailed in Eq. (14)) is such that $\hat{\sigma}^\tau_{\kappa_l} = 0$ for each $\kappa_l \in \mathcal{K}$. In this case, investigation over aspect κ is clearly unable to persuade the Receiver to accept the proposal.

Consider now the case where, for some type τ, the set $\hat{\mathcal{K}}^\tau$ is nonempty so that there is at least one aspect realization $\hat{\kappa}^\tau \in \mathcal{K}$ such that $\eta^\tau_\kappa(\hat{\kappa}^\tau) > 0$. Then, given that the Sender benefits always from the Receiver accepting the proposal—i.e., $\alpha^\tau_\kappa(\kappa_l) > 0$ for each aspect realization $\kappa_l \in \mathcal{K}$—, any solution $\hat{\sigma}^\tau_\kappa$ to problem $[\mathcal{P}]$ must entail $\hat{\sigma}^\tau_{\hat{\kappa}^\tau} = 1$ for each $\hat{\kappa}^\tau \in \hat{\mathcal{K}}^\tau$. Therefore, the Sender is able to persuade the Receiver by recommending acceptance of the proposal with probability one, conditional on the aspect realizations that belong to the agreement set $\hat{\mathcal{K}}^\tau$. Using this optimal choice, and noting that the optimal decision rule will satisfy the incentive-compatibility constraint with equality, it follows that the information design problem $[\mathcal{P}]$ is solved if and only if the following (adjusted) optimization problem is solved:

$$\begin{aligned} \max_{\{\hat{\sigma}^\tau_{\kappa^\tau}\ \mid\ \kappa^\tau \in \mathcal{K}\setminus\hat{\mathcal{K}}^\tau\}} \quad & \sum_{\kappa^\tau \in \mathcal{K}\setminus\hat{\mathcal{K}}^\tau} \alpha^\tau_\kappa(\kappa^\tau)\ \hat{\sigma}^\tau_{\kappa^\tau} + \sum_{\hat{\kappa}^\tau \in \hat{\mathcal{K}}^\tau} \alpha^\tau_\kappa(\hat{\kappa}^\tau) \\ \text{s.t.:} \quad & \sum_{\kappa^\tau \in \mathcal{K}\setminus\hat{\mathcal{K}}^\tau} \eta^\tau_\kappa(\kappa^\tau)\ \hat{\sigma}^\tau_{\kappa^\tau} = \sum_{\hat{\kappa}^\tau \in \hat{\mathcal{K}}^\tau} \eta^\tau_\kappa(\hat{\kappa}^\tau). \end{aligned} \tag{15}$$

Now, in order to analyze the acceptance recommendation probabilities $\{\hat{\sigma}^\tau_{\kappa^\tau} \mid \kappa^\tau \in \mathcal{K} \setminus \hat{\mathcal{K}}^\tau\}$ that solve the linear programming problem described by Eq. (15), we need to compare, for each pair of different aspect realizations $\kappa^\tau_l, \kappa^\tau_g \in \mathcal{K} \setminus \hat{\mathcal{K}}^\tau$, the (directional) slopes $\alpha^\tau_\kappa(\kappa^\tau_l)/\alpha^\tau_\kappa(\kappa^\tau_g)$ of the Sender's indifference lines with the respective (directional) slopes $\eta^\tau_\kappa(\kappa^\tau_l)/\eta^\tau_\kappa(\kappa^\tau_g)$ of the required incentive-compatibility condition. This, in turn, translates into comparing the ratios $\alpha^\tau_\kappa(\kappa^\tau_l)/\eta^\tau_\kappa(\kappa^\tau_l)$ and $\alpha^\tau_\kappa(\kappa^\tau_g)/\eta^\tau_\kappa(\kappa^\tau_g)$ for each pair of different aspect realizations $\kappa_l^\tau, \kappa_g^\tau \in \mathcal{K} \setminus \hat{\mathcal{K}}^\tau$. By resorting to these comparisons, it follows from standard results of linear programming that any solution to the problem in Eq. (15) implies that $\hat{\sigma}_{\bar{\kappa}^\tau}^\tau \geq 0$ only for the aspect realization(s) $\bar{\kappa}^\tau \in \mathcal{K} \setminus \hat{\mathcal{K}}^\tau$ associated with the maximal ratio $\alpha_\kappa^\tau(\kappa_l^\tau)/\eta_\kappa^\tau(\kappa_l^\tau)$ across all possible realizations $\kappa_l^\tau \in \mathcal{K} \setminus \hat{\mathcal{K}}^\tau$, whereas $\hat{\sigma}_{\kappa^\tau}^\tau = 0$ for each aspect realization κ^τ such that $\alpha_\kappa^\tau(\kappa^\tau)/\eta_\kappa^\tau(\kappa^\tau) \notin \max_{\kappa_l^\tau \in \mathcal{K}\setminus\hat{\mathcal{K}}^\tau} \quad \alpha_\kappa^\tau(\kappa_l^\tau)/\eta_\kappa^\tau(\kappa_l^\tau)$. Conditional on information design over aspect κ, this result establishes the conditions on the rates of substitution of the Sender's expected utility, relative to the (directional) slopes of the required incentive-compatibility condition for acceptance, that characterize the Sender's optimal information design choice.

Obviously, the set $\mathcal{A}_\kappa^\tau$ is nonempty since $\mathcal{K} \setminus \hat{\mathcal{K}}^\tau$ is a finite set. Suppose that the set $\mathcal{A}_\kappa^\tau \equiv \arg\max_{\kappa_l^\tau \in \mathcal{K}\setminus\hat{\mathcal{K}}^\tau} \quad \alpha_\kappa^\tau(\kappa_l^\tau)/\eta_\kappa^\tau(\kappa_l^\tau)$ is a singleton so that $\mathcal{A}_\kappa^\tau = \{\bar{\kappa}^\tau\}$. Then, it follows from the form of the problem in Eq. (15) that the Sender optimally chooses $\hat{\sigma}_{\bar{\kappa}^\tau}^\tau = \sum_{\hat{\kappa}^\tau \in \hat{\mathcal{K}}^\tau} \eta_\kappa^\tau(\hat{\kappa}^\tau)/\eta_\kappa^\tau(\bar{\kappa}^\tau) \in (0,1)$. If the set $\mathcal{A}_\kappa^\tau$ is not a singleton, then from the form of the problem in Eq. (15) we observe that the Sender optimally chooses any set of conditional probabilities $\{\hat{\sigma}_{\bar{\kappa}^\tau}^\tau \in [0,1] \mid \bar{\kappa}^\tau \in \mathcal{A}_\kappa^\tau\}$ that satisfies the incentive-compatibility condition

$$\sum_{\bar{\kappa}^\tau \in \mathcal{A}_\kappa^\tau} \eta_\kappa^\tau(\bar{\kappa}^\tau)\ \hat{\sigma}_{\bar{\kappa}^\tau}^\tau = \sum_{\hat{\kappa}^\tau \in \hat{\mathcal{K}}^\tau} \eta_\kappa^\tau(\hat{\kappa}^\tau).$$

This describes completely the optimal solution to the incentive-compatible constrained problem that a Sender of type τ faces when he selects aspect κ for information design and decides to fully reveal his information about the pattern of dependence between the two aspects.

We turn now to study the condition that describes the optimal aspect choice of the Sender, provided that he decides to fully reveal his information about the pattern of dependence between the two aspects. Suppose that the type of the Sender is some given $\tau \in \mathcal{T}$. First, note that if, for some of the aspects $\kappa \in \{x, y\}$, the set $\hat{\mathcal{K}}^\tau$ is empty, then such an aspect is not optimally chosen for information design when the Sender has type τ. Then, consider that both sets $\hat{X}^\tau$ and $\hat{Y}^\tau$ are nonempty so that, as shown above, the Sender can use investigation choices over any aspect, either x or y, to persuade the Receiver by recommending acceptance with probability one for some realizations of both aspects. For this case, take a generic aspect $\kappa \in \{x, y\}$ and pick any aspect realization $\bar{\kappa}^\tau \in \mathcal{A}_\kappa^\tau$. Notice that, in the case where the set $\mathcal{A}_\kappa^\tau$ is not a singleton, investigation can recommend acceptance with positive probability *only* upon the realization of the selected aspect $\bar{\kappa}^\tau$ and yet the Sender's optimal expected utility would be identical to the one that he would obtain when investigation recommends acceptance conditional on any subset of aspect realizations from $\mathcal{A}_\kappa^\tau$, with the restriction that the incentive-compatibility constraint is satisfied. Therefore, we can consider without loss of generality the optimal expected utility that the Sender obtains upon investigation over aspect κ by taking the optimal choice $\hat{\sigma}^{\tau}_{\bar{\kappa}^{\tau}} = \sum_{\hat{\kappa}^{\tau}\in\hat{\mathcal{K}}^{\tau}} \eta^{\tau}_{\kappa}(\hat{\kappa}^{\tau})/\eta^{\tau}_{\kappa}(\bar{\kappa}^{\tau})$. Then, by plugging such an optimal choice into the Sender's objective function of his problem described in Eq. (15), it follows that the optimal expected utility that the Sender of type τ obtains upon information design over aspect κ, and conditional on fully revealing his type, equals

$$\sum_{\hat{\kappa}^{\tau}\in\hat{\mathcal{K}}^{\tau}} \left[\frac{\alpha^{\tau}_{\kappa}(\bar{\kappa}^{\tau})}{\eta^{\tau}_{\kappa}(\bar{\kappa}^{\tau})}\eta^{\tau}_{\kappa}(\hat{\kappa}^{\tau}) + \alpha^{\tau}_{\kappa}(\hat{\kappa}^{\tau})\right],$$

where recall that $\bar{\kappa}^{\tau}$ is an arbitrarily chosen aspect realization from the set $\mathcal{A}^{\tau}_{\kappa}$. Therefore, for the case where the Sender can persuade the Receiver by committing investigation over any of the two aspects of uncertainty—because the sets $\hat{X}^{\tau}$ and $\hat{Y}^{\tau}$ are nonempty—, the Sender picks arbitrarily any aspect realization $\bar{\kappa}^{\tau} \in \mathcal{A}^{\tau}_{\kappa}$ for each of the two aspects $\kappa \in \{x, y\}$ and then optimally selects for investigation the aspect(s) κ^* that solve(s) the problem

$$\max_{\kappa\in\{x,y\}} \sum_{\hat{\kappa}^{\tau}\in\hat{\mathcal{K}}^{\tau}} \left[\frac{\alpha^{\tau}_{\kappa}(\bar{\kappa}^{\tau})}{\eta^{\tau}_{\kappa}(\bar{\kappa}^{\tau})}\eta^{\tau}_{\kappa}(\hat{\kappa}^{\tau}) + \alpha^{\tau}_{\kappa}(\hat{\kappa}^{\tau})\right],$$

as stated. ∎

Proof of Theorem 2.
Suppose that the Sender selects an aspect $\kappa \in \{x, y\}$ to commission investigation.

Consider first the case where the Sender chooses to fully reveal his private information about the pattern of dependence ($d = r$). Let me analyze the information design problem described by Eq. (9), provided that the Sender has type $\tau' \in \mathcal{T}$. It follows from Assumption 2 (iv) on the Sender's preferences that his optimal investigation choice must entail $\sigma^{(\tau',r)^*}_{\kappa}(\underline{a} \mid \kappa) = 0$ for each given realization $\kappa \in \mathcal{K}$. Given this, the problem that a Sender of type τ' must solve, upon fully revealing his private information ($d = r$), in order to choose his decision rule $\hat{\sigma}^{\tau'}_{\kappa} = \{\hat{\sigma}^{\tau'}_{\kappa_l} \in [0, 1]\}^{m}_{l=1}$, has then the form:[33]

$$\begin{aligned} &\max_{\{\hat{\sigma}^{\tau'}_{\kappa}\}} \sum_{\theta\in\Theta} \psi_{\tau'}(\theta)\ u_S(\overline{a}, \theta)\ \hat{\sigma}^{\tau'}_{\kappa_l} \\ &\text{s.t.:} \sum_{\theta\in\Theta} \psi_{\tau'}(\theta)\ u_R(\overline{a}, \theta)\ \hat{\sigma}^{\tau'}_{\kappa_l} \geq 0. \end{aligned} \tag{16}$$

[33] Although presented using a slightly different notation, notice that this problem in Eq. (16) coincides with problem $[\mathcal{P}]$, which was stated in Eq. (14).

Consider now the case where the Sender chooses instead to withhold at least some of his private information about the pattern of dependence ($d \neq r$). Then, the problem that a Sender of type τ such that $d(\tau) = \mathcal{T}_\tau \neq \{\tau\}$ (for some non-singleton subset $\mathcal{T}_\tau \subset \mathcal{T}$) must solve—upon withholding (some of) his private information ($d \neq r$)—in order to choose his decision rule, has the form:

$$\begin{aligned} &\max_{\{\sigma_\kappa^{(\tau,d)}\}} \sum_{\theta\in\Theta} \sigma_\kappa^{(\tau,d)}(\overline{a} \mid \kappa)\ \psi_\tau(\theta)\ u_S(\overline{a},\theta) \\ &\text{s.t.:} \quad \sum_{\tau'\in\mathcal{T}_{\text{skep}(\tau)}} \beta_\tau^d(\tau') \sum_{\theta\in\Theta} \sigma_\kappa^{(\tau',d)}(\overline{a} \mid \kappa)\ \psi_{\tau'}(\theta)\ u_R(\overline{a},\theta) \geq 0 \qquad \forall \tau' \in d(\tau) = \mathcal{T}_\tau \end{aligned} \tag{17}$$

Notice first that, as indicated earlier, an informational consistency criterion imposes that the Receiver must not be able to learn anything about the pattern of dependence by using the aspect-restricted decision rule $\sigma_\kappa^{(\tau,d)}$ optimally chosen by the Sender. In other words, it must be the case that $\sigma_\kappa^{(\tau',d)}$ cannot in fact depend on the type $\tau' \in d(\tau) = \mathcal{T}_\tau$ and, therefore, the incentive-compatibility constraint of the problem in Eq. (17) above can be rewritten as:

$$\sum_{\theta\in\Theta} \sigma_\kappa^d(\overline{a} \mid \kappa)\ u_R(\overline{a},\theta)\ \mathbb{E}_{\beta_\tau^d}[\psi_{\tilde{\tau}}(\theta)] \geq 0, \tag{18}$$

where we are setting $\sigma_\kappa^d \equiv \sigma_\kappa^{(\tau',d)}$ for each $\tau' \in d(\tau) = \mathcal{T}_\tau$. Secondly, for such cases where the Sender does not fully reveal his private information about the pattern of dependence ($d \neq r$), it follows from Assumption 3 that the maximally skeptical Receiver updates her priors about the pattern of dependence by assigning positive probability only to the types of the Sender that minimize his ex ante expected utility. Again, recall that Assumption 2 (iv) on the Sender's preferences implies that the Receiver's ex ante expected utility when the Sender chooses $d = r$ and has type $\tau' \in d(\tau) = \mathcal{T}_\tau$ equals

$$\sum_{a\in A}\sum_{\theta\in\Theta} \sigma_\kappa^{(\tau',r)}(a \mid \kappa)\ \psi_{\tau'}(\theta)\ u_R(a,\theta) = \sum_{\theta\in\Theta} \psi_{\tau'}(\theta)\ u_R(\overline{a},\theta)\ \hat{\sigma}_{\kappa_l}^{\tau'}.$$

Since the Sender's aspect-restricted information design problem has a single incentive-compatibility constraint, then it must be the case that, for the optimal information design choice $\{\hat{\sigma}_{\kappa_l}^{\tau}\}_{l=1}^m$, we have that such a condition holds with equality, i.e., $\sum_{\theta\in\Theta} \psi_\tau(\theta)\ u_R(\overline{a},\theta)\ \hat{\sigma}_{\kappa_l}^{\tau} = 0$ for all $\tau' \in d(\tau) = \mathcal{T}_\tau$. It follows that

$$\mathcal{T}_{\text{skep}}(\tau) = \arg\min_{\{\tau'\in\mathcal{T}_\tau\}} \sum_{\theta\in\Theta} \psi_{\tau'}(\theta)\ u_R(\overline{a},\theta)\ \hat{\sigma}_{\kappa_l}^{\tau'} = \mathcal{T}_\tau.$$

Therefore, the skeptical Receiver may place a positive probability $\beta_\tau^d(\tau')$ on any possible type $\tau' \in \mathcal{T}_\tau$. Under Assumption 3, the Bayesian posterior probability that the Receiver assigns then to each type $\tau' \in \mathcal{T}_\tau$ must equal

$$\beta_\tau^d(\tau') = q(\tau') \Big/ \sum_{\tau'' \in \mathcal{T}_\tau} q(\tau'') \in (0,1).$$

Now, take as given the optimal information design choice $\hat{\sigma}_\kappa^{\tau'}$ that the Sender selects by solving his problem in Eq. (16)—i.e., for the case where the Sender chooses $d = r$, conditional on his type being τ'. Then, suppose that, conditional on choosing $d \neq r$, the Sender commissions investigation by selecting, for any type τ such that $d(\tau) = \mathcal{T}_\tau \neq \{\tau\}$ (for some non-singleton subset $\mathcal{T}_\tau \subset \mathcal{T}$), and then, for each $\tau' \in \mathcal{T}_\tau$, a common decision rule $\sigma_\kappa^d(\overline{a} \mid k_l) = \sigma_\kappa^{(\tau',d)}(\overline{a} \mid k_l)$ by setting

$$\sigma_\kappa^d(\overline{a} \mid k_l) = \sum_{\tau' \in \mathcal{T}_\tau} \omega_\tau(\tau') \, \hat{\sigma}_{\kappa_l}^{\tau'}, \tag{19}$$

where $\omega_\tau : \mathcal{T}_\tau \to (0,1)$ is some positive weight function such that $\sum_{\tau' \in \mathcal{T}_\tau} \omega_\tau(\tau') = 1$. Given a type $\tau' \in \mathcal{T}_\tau$, let me multiply the corresponding incentive-compatibility constraint of the information design problem in Eq. (16) by $\omega_\tau(\tau')\beta_\tau^d(\tau')$ and sum across all possible pairs of types $(\tau', \tau') \in \mathcal{T}_\tau \times \mathcal{T}_\tau$. Since each $\omega_\tau(\tau')\beta_\tau^d(\tau') > 0$ and we have that each decision rule $\hat{\sigma}_\kappa^{\tau'}$ solves the corresponding problem described by Eq. (16), it follows that

$$\begin{aligned}
&\sum_{\tau' \in \mathcal{T}_\tau} \sum_{\tau' \in \mathcal{T}_\tau} \omega_\tau(\tau')\beta_\tau^d(\tau') \sum_{\theta \in \Theta} \psi_{\tau'}(\theta) \; u_R(\overline{a}, \theta) \; \hat{\sigma}_{\kappa_l}^{\tau'} \\
&\qquad = \sum_{\theta \in \Theta} \Big[\sum_{\tau' \in \mathcal{T}_\tau} \omega_\tau(\tau') \; \hat{\sigma}_{\kappa_l}^{\tau'} \Big] \; u_R(\overline{a}, \theta) \; \Big[\sum_{\tau' \in \mathcal{T}_\tau} \beta_\tau^d(\tau') \; \psi_{\tau'}(\theta) \Big] \geq 0 \\
&\qquad \Leftrightarrow \sum_{\theta \in \Theta} \sigma_\kappa^d(\overline{a} \mid k_l) \; u_R(\overline{a}, \theta) \; \mathbb{E}_{\beta_\tau^d}[\psi_{\tilde{\tau}}(\theta)] \geq 0.
\end{aligned}$$

Therefore, the chosen decision rule σ_κ^d, constructed as proposed in Eq. (19), satisfies the incentive-compatibility constraint of the Sender's problem in Eq. (17)—using the particular form detailed in Eq. (18). Moreover, since each $\hat{\sigma}_{\kappa_l}^{\tau'}$ solves the respective information design problem in Eq. (16) and each $\omega_\tau(\tau') > 0$, then the chosen decision rule σ_κ^d solves also the Sender's problem described in Eq. (17). These arguments apply to each type τ such that $d(\tau) = \mathcal{T}_\tau \neq \{\tau\}$ (for some non-singleton subset $\mathcal{T}_\tau \subset \mathcal{T}$). Notice then that, for a message strategy $d \neq r$, we have that, for any $\tau \in \mathcal{T}$ such that $d(\tau) = \mathcal{T}_\tau \neq \{\tau\}$ (for some non-singleton subset $\mathcal{T}_\tau \subset \mathcal{T}$), each type $\tau' \in \mathcal{T}_\tau$ receives an (incentive-compatible constrained) optimal expected utility equal to

$$\tilde{U}_R^\tau \equiv \sum_{\tau' \in \mathcal{T}_\tau} \omega_\tau(\tau') \sum_{\theta \in \Theta} \psi_{\tau'}(\theta)\ u_S(\overline{a}, \theta)\ \ \hat{\sigma}_{\kappa_l}^{\tau'}.$$

Therefore, under Assumption 2 (v), it must be the case that there is some type

$$\bar{\tau} \equiv \arg\max_{\tau' \in \mathcal{T}_\tau} \sum_{\theta \in \Theta} \psi_{\tau'}(\theta)\ u_S(\overline{a}, \theta)\ \hat{\sigma}_{\kappa_l}^{\tau'} \in d(\tau) = \mathcal{T}_\tau$$

such that $\bar{\tau} > \tilde{U}_R^\tau$. Thus, under Assumption 2 (v), type $\bar{\tau}$ has strict incentives to deviate to revealing his private information according to a disclosure strategy d' such that $d'(\bar{\tau}) = \{\bar{\tau}\}$. By applying recursively these arguments, it follows that each type $\tau' \in d(\tau) = \mathcal{T}_\tau$ has strict incentives to deviate and to separate from the rest of pooled types. Notice that all these arguments apply analogously to any message strategy d such that some types pool by sending a common message. Crucially, the arguments presented do not depend on whether the information revealed by each $d(\tau)$ is verifiable or pure cheap talk. Therefore, in each equilibrium ϕ^* of the proposed information disclosure game, the Sender optimally chooses to fully reveal his private information about the pattern of dependence, $d^* = r$, as stated. ∎

Bibliography

Acemoglu, D., V. Chernozhukov, and M. Yildiz (2016): "Fragility of Asymptotic Agreement under Bayesian Learning," *Theoretical Economics*, 11, 187–225.

Ambrus, A. and S. Takahashi (2008): "Multi-Sender Cheap Talk with Restricted State Spaces," *Theoretical Economics*, 3, 1–27.

Baron, D. P. (2006): "Persistent Media Bias," *Journal of Public Economics*, 90, 1–36.

Battaglini, M. (2002): "Multiple Referrals and Multidimensional Cheap Talk," *Econometrica*, 70, 1379–1401.

Bergemann, D. and S. Morris (2013): "Robust Predictions in Games with Incomplete Information," *Econometrica*, 81, 1251–1308.

——— (2016): "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games," *Theoretical Economics*, 11.

——— (2019): "Information Design: A Unified Perspective," *Journal of Economic Literature*, 57, 1–57.

Besley, T. and A. Prat (2006): "Handcuffs for the Grabbing Hand? Media Capture and Government Accountability," *American Economic Review*, 96, 720–736.

Bester, H. and R. Strausz (2001): "Contracting with Imperfect Commitment and the Revelation Principle: The Single Agent Case," *Econometrica*, 69, 1077–1098.

Bewley, T. F. (2002): "Knightian Decision Theory: Part I," *Decisions in Economics and Finance*, 25, 79–110.

Chakraborty, A. and R. Harbaugh (2007): "Comparative Cheap Talk," *Journal of Economic Theory*, 132, 70–94.

——— (2010): "Persuasion by Cheap Talk," *American Economic Review*, 100, 2361–2382.

Che, Y.-K., W. Dessein, and N. Kartik (2013): "Pandering to Persuade," *American Economic Review*, 103, 47–79.

Che, Y.-K. and N. Kartik (2009): "Opinions as Incentives," *Journal of Political Economy*, 117, 815–860.

Crawford, V. P. and J. Sobel (1982): "Strategic Information Transmission," *Econometrica*, 5, 1431–1451.

Deimen, I. and D. Szalay (2019): "Delegated Expertise, Authority, and Communication," *American Economic Review*, 109, 1349–1374.

Frankel, A. and N. Kartik (2019): "Muddled Information," *Journal of Political Economy*, 127.

Frechette, G. R., A. Lizzeri, and J. Perego (2019): "Rules and Commitment in Communication: An Experimental Analysis," Tech. rep., New York University.

Gentzkow, M. and E. Kamenica (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.

——— (2017c): "Disclosure of Endogenous Information," *Economic Theory Bulletin*, 5, 47–56.

Gentzkow, M. and J. M. Shapiro (2006): "Media Bias and Reputation," *Journal of Political Economy*, 114, 280–316.

——— (2010): "What Drives Media Slant? Evidence from US Daily Newspapers," *Econometrica*, 78, 35–71.

Gilboa, I. and D. Schmeidler (1989): "MaxMin Expected Utility with Non-Unique Prior," *Journal of Mathematical Economics*, 18, 141–153.

Glazer, J. and A. Rubinstein (2004): "On Optimal Rules of Persuasion," *Econometrica*, 72, 1715–1736.

Green, J. R. and N. L. Stokey (1980): "A Two-Person Game of Information Transmission," Tech. rep., Harvard University.

Grossman, S. J. (1981): "The Informational Role of Warranties and Private Disclosure about Product Quality," *Journal of Law and Economics*, 24, 461–483.

Hayakawa, S. I. (1940): *Language in Thought and Action*, New York: Harcourt, Brace and Company.

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

Kartik, N., F. X. Lee, and W. Suen (2017): "Investment in Concealable Information by Biased Experts," *RAND Journal of Economics*, 48, 24–43.

Klibanoff, P., M. Marinacci, and S. Mukerji (2005): "A Smooth Model of Decision Making under Ambiguity," *Econometrica*, 73, 1849–1892.

Krishna, V. and J. Morgan (2001): "Asymmetric Information and Legislative Rules: Some Admendments," *American Political Science Review*, 95, 435–452.

Levy, G. and R. Razin (2007): "On the Limits of Communication in Multidimensional Cheap Talk: A Comment," *Econometrica*, 75, 885–803.

Lipnowski, E. and D. Ravid (2019): "Cheap Talk with Transparent Motives," Tech. rep., University of Chicago.

Lipnowski, E., D. Ravid, and D. Shishkin (2019): "Persuasion via Weak Institutions," Tech. rep., University of Chicago.

Maccheroni, F., M. Marinacci, and A. Rustichini (2006a): "Ambiguity Aversion, Robustness, and the Variational Representation of Preferences," *Econometrica*, 74, 1447–1498.

——— (2006b): "Dynamic Variational Preferences," *Journal of Economic Theory*, 128, 4–44.

Milgrom, P. and J. Roberts (1986): "Relying on the Information of Interested Parties," *RAND Journal of Economics*, 17, 18–32.

Milgrom, P. R. (1981): "Good News and Bad News: Representation Theorems and Applications," *The Bell Journal of Economics*, 12, 380–391.

MIN, D. (2017): "Bayesian Persuasion under Partial Commitment," Tech. rep., University of Arizona.

MULLAINATHAN, S. AND A. SHLEIFER (2005): "The Market for News," *American Economic Review*, 95, 1031–1053.

NGUYEN, A. AND T. Y. TAN (2019): "Bayesian Persuasion with Costly Messages," Tech. rep., Carnegie Mellon University.

RAYO, L. AND I. R. SEGAL (2010): "Optimal Information Disclosure," *Journal of Political Economy*, 118, 949–987.

SEO, K. (2009): "Ambiguity and Second-Order Belief," *Econometrica*, 77, 1575–1605.

SHIN, H. S. (1994): "News Management and the Value of Firms," *The RAND Journal of Economics*, 25, 58–71.

——— (2003): "Disclosures and Asset Returns," *Econometrica*, 71, 105–133.

TANEVA, I. (2018): "Information Design," Tech. rep., The University of Edinburgh.