THEMA

théorie économique, modélisation et applications

THEMA Working Paper n°2022-13
CY Cergy Paris Université, France

The perils of a coherent narrative

Alessandro Ispano

DECEMBER 2023

The perils of a coherent narrative

Alessandro Ispano

December 2023

Abstract

A persuader influences a decision-maker by providing a model for interpreting some upcoming news. The decision-maker adopts the model if it does not distort the marginal distribution of news. Both parties can benefit if the persuader can provide news contingent, overall incoherent, models, privately learn the truth, or design the process of news arrival.

Keywords: interpretation, consistency, misspecification, manipulation, model, cognition
JEL classifications: D82, D83, D90

1 Introduction

Recent literature suggests that persuasion often occurs through the provision of an interpretation, or *narrative*, for commonly available information (Eliaz and Spiegler, 2020; Eliaz et al., 2021; Schwartzstein and Sunderam, 2021), rather than through the strategic revelation of private information (Milgrom, 1981; Crawford and Sobel, 1982) or the design of available information (Kamenica and Gentzkow, 2011). Examples range from public policy, e.g., the government claiming that health indicators justify a lockdown during a pandemic, to finance, e.g., an advisor suggesting that stock market returns are favorable for his client's investment, and research, e.g., an empirical study documenting a significant treatment effect in the data. This short paper uses a canonical persuasion game and a standard assumption on admissible narratives to identify some counterintuitive welfare implications that can arise when the persuader must propose a narrative before knowing the actual information that will become available.

Alessandro Ispano: CY Cergy Paris Université, CNRS, THEMA, F-95000 Cergy, France, alessandro.ispano@gmail.com. I am thankful to Kai Barron, Frédéric Koessler, Andrei Matveenko, Alfonso Montes, Egor Starkov, Peter Vida as well as seminar participants at THEMA, Paris II and Paris I for useful comments. Financial support from Labex MME-DII and ANR grant 19-CE26-0010-03 is gratefully acknowledged.

In our framework, a persuader (he) influences a decision-maker (she) by providing a narrative for interpreting some upcoming, uncertain, news. The decision-maker adopts any proposed narrative that leaves the marginal distribution of the news undistorted relative to the true model. After the news is realized, the decision-maker updates the prior accordingly and chooses an action. A natural interpretation is that the decision-maker has knowledge of the frequency of states, as specified by the prior distribution, and of each news, as specified by the true news generating process. However, she is uncertain about the correlation between the two, possibly due to the unavailability of detailed historical data. She is therefore naively willing to accept any interpretation provided by the "expert" that is consistent with her knowledge.

We compare the set of beliefs that the persuader can induce under this ex-ante timing with the ex-post timing in which the persuader proposes a narrative only once she knows the news. Given that the persuader has state-independent preferences, the ex-post timing is equivalent to a fictitious ex-ante timing scenario in which the persuader can propose, and the decision-maker adopt, narratives that are not *coherent*, i.e. such that the likelihoods of news conditional on states do not form a proper probability distribution.

The coherence requirement implied by the ex-ante timing restricts the set of beliefs that the persuader can induce and thus always harms him in equilibrium. Yet, it sometimes harms also the decision-maker. Indeed, when unfavorable news is more likely ex-ante, the persuader's narrative may interpret it as favorable to induce the decision-maker to take his preferred action more often. Since not all news can be good, the narrative also necessarily interprets favorable news as unfavorable, while an incoherent narrative could at least continue to present favorable news as such. Moreover, because of the ex-ante timing, both players may benefit when the persuader knows the true state of the world. Indeed, when the state is good, good news may become more likely, and the persuader's interpretation then preserves the true meaning of the news. Likewise, both players may be better off when it is the persuader who designs the process of news arrival, rather than it being exogenously given. He would then have no incentive to propose a false narrative in the first place. The concluding section discusses more general formulations and implications of these results and alternative assumptions on the adoption of narratives.

Relation to the literature The recent, growing, literature on model persuasion features two main alternative assumptions about the boundedly rational process of narrative adoption. One branch posits that narratives must be consistent with aspects of the true data generating process, in particular, following the causal misspecification formalization of Spiegler (2016), that the marginal distribution of each variable must be correct (Eliaz and Spiegler, 2020; Eliaz et al., 2021). Another posits that agents adopt narratives that are ex-post more plausible given the data (Schwartzstein and Sunderam, 2021).[1] This paper follows the former approach, albeit in a setting with only one variable in addition to the payoff-relevant one. To the best of our knowledge, the welfare implications it derives in this simple setting by focusing on the ex-ante timing are novel. Two closely related, contemporary papers are Espitia (2023) and Aina (2022). While the focus of Espitia (2023) is rather different, i.e., a principal must choose among agents based on their beliefs about the underlying model, he allows for analogous misspecifications as in this paper. He shows how the ideal bias in the agent's beliefs depends on features of the bias in her preferences. While he restricts his attention to the ex-ante timing, we conjecture that, for the class of state-dependent preferences he considers, the ex-post timing is never beneficial to the agent. Aina (2022) adopts the plausibility assumption of Schwartzstein and Sunderam (2021), in which the persuader proposes an interpretation only after the news is realized, but considers the ex-ante timing. She shows how the persuader can benefit from proposing multiple interpretations, each tailored to different news. As a result, the persuader can induce beliefs that are incoherent from an ex-ante perspective. She also shows that the persuader is not harmed by the ex-ante timing if the decision-maker is initially endowed with a model. Finally, some results of this paper are related to work that, while using the plausibility assumption, combines information design and model persuasion (Ichihashi and Meng, 2021; Jain, 2023).

2 The game

We consider the following persuasion game taking place between a sender (S, he) and a receiver (R, she). The state of the world ω can be good ($\omega = G$) or bad ($\omega = B$) and R has to decide whether to invest ($a = 1$) or not ($a = 0$). S wants to persuade R to invest regardless of the state, i.e. his payoff is $U_S(a) = a$. Instead, R wants to invest only if the state is good, i.e. her payoff $U_R(a, \omega)$ is 1 if $a = 1$ and $\omega = G$, $c/(1-c)$ if $a = 0$ and $\omega = 0$, and 0 otherwise, where $c \in (0, 1)$ is a known parameter. We identify a distribution over states with the probability that the state is good, i.e. $\mu = \mathbb{P}(\omega = G)$, and $\mu_0 \in (0, c)$ represents the prior probability. Thus R finds it optimal to invest only if her belief μ is at least c.

Before choosing her action, R observes the realization of an exogenously given signal π characterized by realization space $X = \{b, g\}$ and likelihoods $\pi(b|B)$ and $\pi(g|G)$, where $\pi(x|\omega)$

[1] See Kendall and Charles (2022) and Barron and Tilman (2022) for experimental evidence on the former and the latter, respectively.

denotes the probability of observing the realization x when the state is ω. Without loss of generality, we assume that $\pi(g|G) \geq 1-\pi(b|B)$, so that b and g represent the "bad" and "good" realization, respectively.

At the initial stage S, knowing the true model π that governs signal realizations, but not the actual realization, may persuade R to believe in a different model. Formally, a model m specifies for each state ω and realization $x \in X$ a probability $\pi_m(x|\omega) \in [0,1]$ with the property, which we call *coherence*, that

$$\sum_{x \in X} \pi_m(x|\omega) = 1 \quad \text{for each } \omega. \tag{1}$$

We denote the space of all possible models by $\mathcal{M}$ and we drop the subscript m when an expression is evaluated according to the true model. R adopts any proposed model $m \in \mathcal{M}$ that is *compatible* with the true model, i.e., such that for each $x \in X$

$$\mathbb{P}_m(x) = \mathbb{P}(x), \tag{2}$$

where $\mathbb{P}_m(x)$ and $\mathbb{P}(x)$ denote the probability of observing the realization x under the model m and the true model, respectively. If S does not propose a model or his proposition does not satisfy this property, R sticks to the true model or, equivalently, to the prior.[2]

We will compare this *ex-ante* timing of S's model proposal with the *ex-post* timing in which S proposes a model only after observing the signal realization. The exact timing of R's model adoption, i.e., whether before or after observing the signal realization, does not matter. Throughout, an equilibrium refers to a pair of strategies of S and R such that S's model proposal is optimal given R's decisions and R's decisions are optimal with respect to the (possibly false) beliefs induced by Bayes' rule under the model R adopts as specified above. We also solve any multiplicity of equilibrium outcomes that may exist for nongeneric parameters by focusing on S's preferred equilibria and, among them, on R's preferred equilibria.

[2] S can always induce R to adopt the true model or one according to which the signal is uninformative, i.e. $\pi_m(x|\omega) = \mathbb{P}(x)$ for any $x \in X$, which is compatible and coherent by construction.

3 Analysis

3.1 The role of coherence in limiting manipulation

Let $\mu_m(x)$ denote the belief induced by model m at realization x and say that a belief pair $\{\mu_m(b), \mu_m(g)\}$ is implementable if there exists a model inducing such beliefs that R will accept.

Observation 1. *The set of implementable beliefs under the ex-ante timing is*

$$\{\mu_m(b), \mu_m(g)\} : \mu_m(b) = \frac{\mu_0}{\mathbb{P}(b)} - \frac{\mathbb{P}(g)}{\mathbb{P}(b)}\mu_m(g) \tag{3}$$

with $max\left\{0, 1 - \frac{1-\mu_0}{\mathbb{P}(x)}\right\} \leq \mu_m(x) \leq min\left\{\frac{\mu_0}{\mathbb{P}(x)}, 1\right\}$ *for each* $x \in \{b, g\}$.

Proof. A model m is accepted if and only if for the realization $x = g$ we have

$$\mathbb{P}(g) = \mu_0\pi_m(g|G) + (1 - \pi_m(b|B))(1 - \mu_0). \tag{4}$$

This follows from equation (2), the definition of $\mathbb{P}_m(g)$ and the fact that if equation (4) holds, i.e. if $\mathbb{P}_m(g) = \mathbb{P}(g)$, it will also be the case that $\mathbb{P}_m(b) = \mathbb{P}(b)$, since by construction $\mathbb{P}_m(g) + \mathbb{P}_m(b) = 1$. Equation (3) is obtained by rearranging terms using that $\mu_0\pi_m(g|G)/\mathbb{P}(g) = \mu_m(g)$ and $(1 - \mu_0)\pi_m(b|B)/\mathbb{P}(b) = 1 - \mu_m(b)$. The bounds, and hence the result, follow from the constraint that $\mu_m(x) \in [0, 1]$ or, equivalently, that $\pi_m(x|G) \in [0, 1]$ for each x. □

Thus, in the space of posterior beliefs, the set of implementable beliefs under the ex-ante timing can be represented as a decreasing line passing through the prior beliefs and the posterior beliefs induced by the true model (see figure 1).[3] It follows from this observation that an accepted model cannot move both beliefs in the same direction. In fact, the distribution of posteriors must be Bayes' plausible, i.e., the expected induced posterior probability is equal to the prior.[4]

Suppose instead that S can propose the model ex-post, after observing the realization x.

Observation 2. *The set of implementable beliefs under the ex-post timing is*

$$\{\mu_m(b), \mu_m(g)\} : max\left\{0, 1 - \frac{1-\mu_0}{\mathbb{P}(x)}\right\} \leq \mu_m(x) \leq min\left\{\frac{\mu_0}{\mathbb{P}(x)}, 1\right\} \quad \textit{for each} \quad x \in \{b, g\}. \tag{5}$$

Proof. For realization x, model m is accepted if and only if

$$\mathbb{P}(x) = \mu_0\pi_m(x|G) + (\pi_m(x|B))(1 - \mu_0), \tag{6}$$

[3]This graphical representation of implementable beliefs is taken from Aina (2022). Instead, Espitia (2023) provides a characterization in terms of implementable joint distributions of signal realizations and states.

[4]Indeed, in the terminology of Spiegler (2020), the causal representation of R is (trivially) a perfect graph.

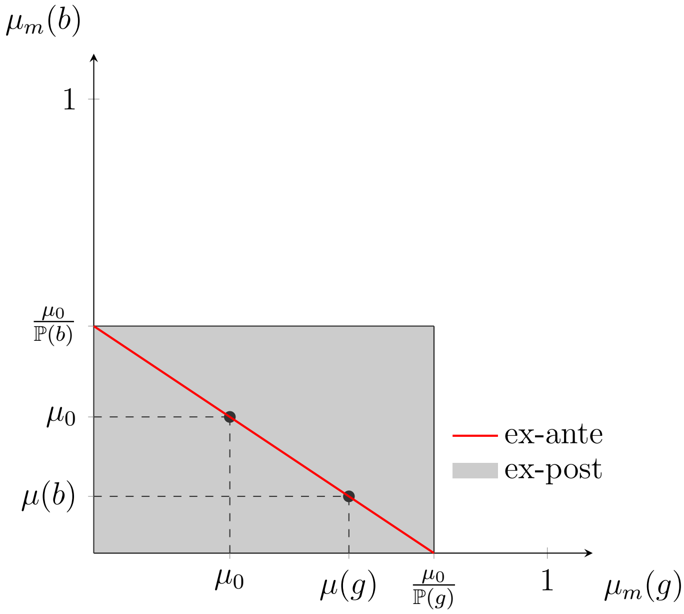

Figure 1 Implementable beliefs ex-ante and ex-post ($\mu = 0.3$, $\pi(g|G) = \pi(b|B) = 0.75$)

where again, the only if part follows from equation (2) and the if part from the fact that equation (6) then also holds for the other realization $\tilde{x}$ (with $\pi_m(\tilde{x}|\omega) = 1 - \pi_m(x|\omega)$ for each ω). The bounds for $\mu_m(x)$ again follow from its definition and the constraint that $\pi_m(x|G) \in [0,1]$ and $\pi_m(x|B) \in [0,1]$. □

Thus, the set of implementable beliefs under the ex-post timing can be represented as a rectangle whose diagonal is the set of implementable beliefs under the ex-ante timing (see again figure 1).

Observation 3. *Consider the ex-ante timing but suppose that S could propose and R could accept incoherent models, i.e., models for which equation* (1) *does not hold. Then the set of implementable beliefs is the same as under the ex-post timing.*

Proof. This follows from the fact that, without the coherence requirement, the choice of $\pi_m(x|\omega)$ for each realization x must simply satisfy equation (6), and is therefore independent of the choice of $\pi_m(\tilde{x}|\omega)$ for the other realization $\tilde{x}$. □

Since S has state-independent preferences, we can henceforth refer to the ex-post timing or to ex-ante timing without coherence interchangeably. Indeed, since knowledge of x does not affect S' optimal model proposition, the outcome of the game will be the same in both scenarios.

3.2 The perils of coherence

Given the characterization of implementable beliefs in the previous section, it is straightforward to deduce the outcome of the game with and without the coherence requirement. Without coherence, S will be able to induce R to invest for realization $x \in \{b, g\}$ if and only if

$$c \leq \frac{\mu_0}{\mathbb{P}(x)}, \tag{7}$$

namely, as noted by Schwartzstein and Sunderam (2021), if the prior μ_0 is sufficiently close to the target belief c or the realization x is sufficiently unlikely. And if such a condition holds for both realizations, R will always invest. With coherence, instead, so that S cannot induce R to always invest, S would rather induce R to invest for the more likely realization if possible, i.e. if it is such that equation (7) holds. These observations imply the following proposition.

Proposition 1 (Coherence and receiver welfare). *If $\mathbb{P}(b) > \mathbb{P}(g)$, $c \leq \frac{\mu_0}{\mathbb{P}(b)}$ and $\mu(g) > c$ then R is strictly worse off under the ex-ante timing than under the ex-post timing. In all other cases, R is weakly better off, and sometimes strictly, under the ex-ante timing.*

Proof. The first inequality states that under the ex-ante timing S strictly prefers to induce R to invest for the bad realization, the second that it is possible to do so, and the third that, given the true model, R should invest upon good news. Then, in equilibrium, she will do so under the ex-post timing, but not under the ex-ante timing. In all other cases, R weakly benefits because either her decisions will be the same under the ex-ante and the ex-post timing, or she will invest more often under the ex-post timing when unwarranted (e.g., when $\mathbb{P}(b) > \mathbb{P}(g)$ and $c \leq \frac{\mu_0}{\mathbb{P}(b)}$ but $\mu(g) < c$). □

As an illustrative example, we will consider the case $\pi(b|B) = \pi(g|G) = \rho > 1/2$ and $c = 1/2$, i.e., a symmetric signal with precision ρ and symmetric gains for R when her action matches the state. Since $\mathbb{P}(b)$ is increasing in ρ and decreasing in μ_0, S will be able to induce R to invest for the bad, more likely, realization if μ_0 is sufficiently large or ρ is sufficiently small, i.e., if $\rho \leq \bar{\rho}(\mu_0) \equiv \frac{\mu_0}{1-2\mu_0}$, where cutoff $\bar{\rho}$ is the solution to $1/2 = \mu_0/\mathbb{P}(b)$ with respect to ρ. Cutoff $\bar{\rho}$ is increasing in μ_0 and the condition is always satisfied if $\mu_0 \geq 1/3$ and always violated if $\mu_0 < 1/4$. If the condition is violated, S can at least induce R to invest for the good, less likely, realization if $\rho \geq \underline{\rho}(\mu_0) \equiv \frac{1-3\mu_0}{1-2\mu_0}$, where cutoff $\underline{\rho}$ is the solution to $1/2 = \mu_0/\mathbb{P}(g)$ with respect to ρ. Cutoff $\underline{\rho}$ is decreasing in μ_0 and the condition is always satisfied when $\mu_0 \geq 1/4$.

Figure 2 illustrates the three regions that partition the parameter space based on R's investment decision under S's optimal coherent model proposition. The dashed line, of equation $1 - \mu_0$, represents the precision level of the signal above which R should invest upon good news under the true model, while investing upon bad news is obviously never warranted. The outcome of the game under the ex-ante and ex-post timing differs in the region C, where R would always

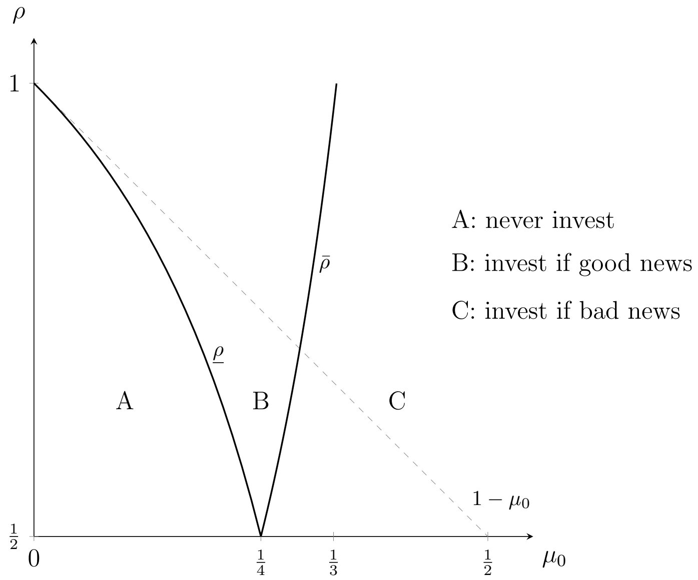

Figure 2 Equilibrium behavior under the ex-ante timing

invest in the latter case. Then, above the dashed line, R would be better off if S's model did not need to be coherent, since in that case she would at least correctly invest upon the favorable realization. Conversely, the coherence requirement helps R when $\rho < 1 - \mu_0$, since she will at least not invest incorrectly upon the favorable realization. Figure 3 shows the associated equilibrium payoff of R (left panel) and S (right panel) as a function of the precision of the signal ρ for three different values of the prior, namely, highly unfavorable ($\mu_0 < 1/4$, figures 3a and 3b), unfavorable ($\mu_0 \in (1/4, 1/3)$, figures 3c and 3d), and moderately unfavorable ($\mu_0 \in (1/3, 1/2)$, figures 3e and 3f). In particular, due to the coherence requirement, the payoffs and R and S can be minimized and maximized, respectively, when the signal is fully informative.

3.3 Other welfare implications of the ex-ante timing

3.3.1 Benefits from an informed sender

Suppose that S can observe the state, but still not the signal realization, when proposing a model. Since, conditional on the state, the realization that is more likely to be observed may differ relative to the prior, S's optimal model proposal may also change. As a result, not only S but also R may benefit.

Proposition 2. *If $\mathbb{P}(b) > \mathbb{P}(g)$, $c \leq \frac{\mu_0}{\mathbb{P}(b)}$ and $\pi(g|G) > \pi(b|G)$, R is strictly better off when S knows the state.*

Proof. When $\omega = B$, $\mathbb{P}(b|B) = \pi(b|B) \geq \mathbb{P}(b) = \mu_0(1 - \pi(g|G)) + (1 - \mu_0)\pi(b|B)$, since $\pi(b|B) \geq (1 - \pi(g|G)$ by assumption. Thus, in this case, S's optimal model proposition is

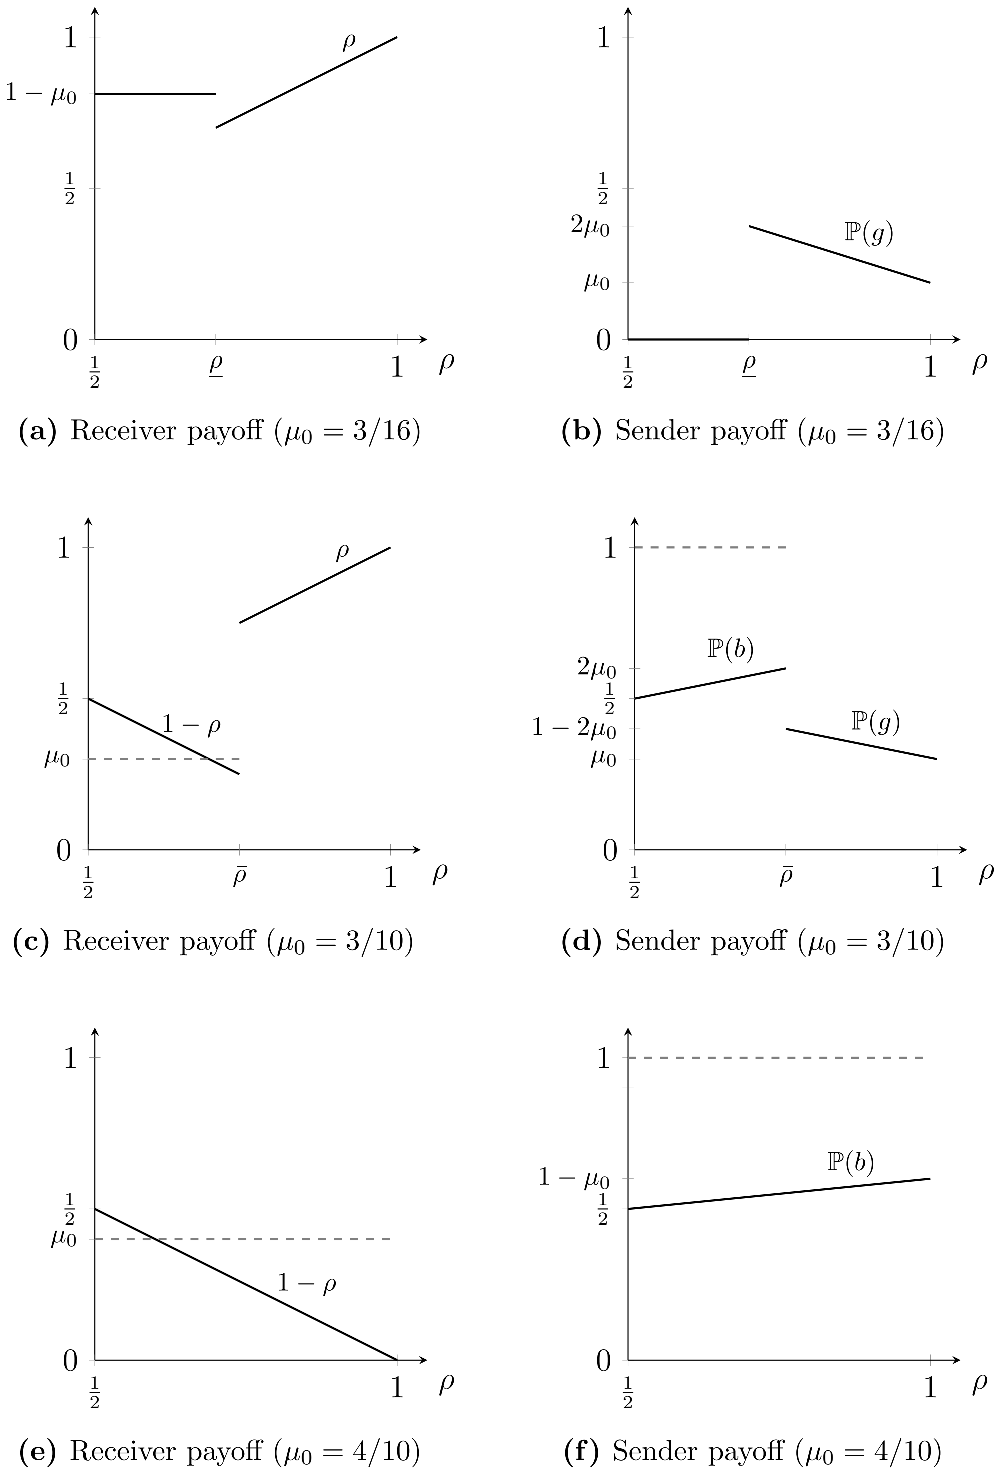

(a) Receiver payoff ($\mu_0 = 3/16$) **(b)** Sender payoff ($\mu_0 = 3/16$)

(c) Receiver payoff ($\mu_0 = 3/10$) **(d)** Sender payoff ($\mu_0 = 3/10$)

(e) Receiver payoff ($\mu_0 = 4/10$) **(f)** Sender payoff ($\mu_0 = 4/10$)

Figure 3 Equilibrium payoffs under the ex-ante timing as a function of signal precision (the dashed lines represent payoffs under the ex-post timing, whenever they differ) unaffected relative to when S does not know ω, and so is R's expected payoff conditional on $\omega = B$. When $\omega = G$, instead, if $\pi(g|G) > \pi(b|G)$, S now finds it optimal to propose a model that induces R to invest for good news, and R's expected payoff conditional on $\omega = G$ increases from $\pi(b|G)$ to $\pi(g|G)$. $\square$

Note that if learning the state entailed a cost to S, he would never acquire information under the ex-post timing. Under the ex-ante timing, S will do so when his optimal model proposition varies with the state, as in the case of proposition 2, and the cost is sufficiently small. Of course, there are also cases in which S's information acquisition harms R by inducing her to invest upon bad news when the state is bad.

3.3.2 Benefits from the sender designing the true model

Now suppose that it is S who designs the true model in the initial stage, i.e. who chooses $\pi(b|B)$ and $\pi(g|G)$. S necessarily benefits from being able to control the process of news arrival. She can then attain her payoff under the classical Bayesian persuasion problem in which S designs the true model and R rationally updates accordingly (Kamenica and Gentzkow, 2011). Again, R may also benefit.

Proposition 3. *If $\mathbb{P}(b) > \mathbb{P}(g)$ and either $c \leq \frac{\mu_0}{\mathbb{P}(b)}$ or $\frac{\mu_0}{\mathbb{P}(g)} \leq c < \frac{\mu_0}{\mathbb{P}(b)}$ and $\mu(g) < c$, or if $\mathbb{P}(g) \geq \mathbb{P}(b)$ and either $\frac{\mu_0}{\mathbb{P}(g)} < c \leq \frac{\mu_0}{\mathbb{P}(b)}$ or $c \leq \frac{\mu_0}{\mathbb{P}(g)}$ and $\mu(g) < c$, R is strictly better off when S designs the true model.*

Proof. The set of belief pairs S can implement when designing the true model is the union of those at observation 1 as $\mathbb{P}(g)$ varies from 0 to 1 (in the two limit cases such a set is the singleton (μ_0, μ_0)). Graphically, it consists of all points of the unit square touched by the line in figure 1 as it rotates around the point (μ_0, μ_0) with a slope varying from 0 to minus infinity. However, to find S's optimum, given that he has state-independent preferences, we can without loss of generality restrict our attention to the case in which he truthfully proposes the model he designs. Indeed, any pair of beliefs that can be induced given the designed model as per observation 1 can be equivalently induced by choosing the proposed model as the true one. This choice gives S the same expected payoff since the two realizations will still have the same probabilities. And even if there are multiple true models that are optimal for S, this choice necessarily yields R's preferred equilibrium, since she updates according to the correct model. Thus, using the results of Kamenica and Gentzkow (2011), the optimal true model of S induces the beliefs $\mu_m(b) = 0$ and $\mu_m(g) = c$, with probabilities $1 - \mu_0/c$ and μ_0/c, respectively, i.e., $\pi(g|G) = 1$ and $\pi(b|B) = \frac{(1-\frac{\mu_0}{c})}{1-\mu_0}$. The expected payoff of R is then $(1 - \mu_0)\frac{c}{1-c}$, i.e., the same as she would get when sticking to the prior, in which case she will never invest. It follows that when the true model is given exogenously, R's expected payoff is strictly lower whenever she makes an unjustified decision to invest given the true model, i.e., for bad news or for insufficiently good news. These are the cases covered by the proposition. □

In this setting, S's choice of an ex-ante model proposition for an exogenously given model essentially amounts to a constrained problem of Bayesian persuasion in which realizations must have fixed probabilities. Therefore, S cannot benefit from the bias in R's beliefs relative to the full-rationality benchmark with unrestricted information design.[5] Thus, there always is an optimal model S would design which is manipulation proof, i.e., such that S has no incentive to propose a false one. R is then better off than when the model is given exogenously if in that case she would make some irrational decisions.

4 Discussion

The common theme of the results in this paper is that a decision-maker who is boundedly rational in her model adoption process may be better off if the persuader is given more discretion. This discretion can take the form of knowing the actual news before proposing a model (proposition 1), of knowing the true (proposition 2), or of being able to control the process of news arrival (proposition 3). Taken literally, these results imply, for example, that, contrary to conventional wisdom, it may be better for a credulous decision-maker to let the persuader come up with an interpretation for the observed facts only once these have realized. And if the decision-maker has some control over these decisions, she may benefit from letting the persuader observe the facts or discover the truth. Likewise, the decision-maker may benefit from letting the persuader suggest what the relevant facts to pay attention to in the future are in the first place. If the ex-post timing is interpreted as standing for a decision-maker who is persuaded even by incoherent narratives as per observation 3, the paper illustrates how an agent's welfare may not be monotonic in her degree of rationality.

More fundamentally, the paper demonstrates how model persuasion can lead to narratives that are sufficiently strange to sound unrealistic even to a boundedly rational agent. These findings suggest the need for a more comprehensive understanding and more robust formaliza-

[5] This result is not generally true. For example, consider the example of figure 2, but suppose that S's payoff changes to $U_S(a, \omega) = \mathbb{1}_{(a=1\&\omega=B)\text{ or }(a=0\&\omega=G)}$, i.e., S wants to induce R to always take the "wrong" action. Consider again the portion or region C above the $1-\mu_0$ line. Since S induces R to take the wrong action according to the true posterior, S's optimal model proposition is unaffected. However, S's payoff increases from $\mathbb{P}(b)$ to ρ. S's payoff is now higher than his Bayesian persuasion one (μ_0).

tion of how agents with limited rationality respond to model persuasion. One potential solution is to adopt a more structured and context-dependent approach to narrative adoption. Indeed, under the standard compatibility assumption used, the true data generating process affects the decision-maker's model adoption decision only by determining the probabilities that the persuader's proposed interpretation must attach to each news. This assumption may be reasonable in contexts that are unfamiliar to the decision-maker, such as an inexperienced investor dealing with financial data, the general population being exposed to health indicators during an unprecedented pandemic, or a reader of a complex empirical study who is not trained in econometrics. However, in other situations, the decision-maker's knowledge may impose additional constraints. For instance, if it is evident that certain news is more favorable than others, the decision-maker may only consider models that preserve this property. More generally, the decision-maker may be unwilling to consider models that imply an excessive distance between the true and induced joint distributions over states and news. In the game under consideration, the outcome with and without coherence may then be similar.

In other situations, the decision-maker may have a biased initial view of the distribution of news, and possibly also of states, i.e., a false prior. In this case, to be accepted, an interpretation may have to be compatible with respect to this subjective data generating process rather than the true one. Interestingly, since biased beliefs are sometimes more immune to manipulation, the decision-maker may be better off. Matters become more complicated when the set of news the decision-maker considers possible does not coincide with the true one, which may require entering the realm of updating upon an unforeseen contingency (see Galperti (2019)).

References

Aina, Chiara, "Tailored Stories," *Working paper*, 2022.

Barron, Kai and Fries Tilman, "Narrative Persuasion," *Working paper*, 2022.

Crawford, Vincent P. and Joel Sobel, "Strategic Information Transmission," *Econometrica*, 1982, *50* (6), pp. 1431–1451.

Eliaz, Kfir and Ran Spiegler, "A model of competing narratives," *American Economic Review*, 2020, *110* (12), 3786–3816.

_ , _ , and Yair Weiss, "Cheating with models," *American Economic Review: Insights*, 2021, *3* (4), 417–34.

Espitia, Andrés, "Confidence and Organizations," *Working paper*, 2023.

Galperti, Simone, "Persuasion: The art of changing worldviews," *American Economic Review*, 2019, *109* (3), 996–1031.

Ichihashi, Shota and Delong Meng, "The Design and Interpretation of Information," *Working paper*, 2021.

Jain, Atulya, "Informing agents amidst biased narratives," *Working paper*, 2023.

Kamenica, Emir and Matthew Gentzkow, "Bayesian persuasion," *American Economic Review*, 2011, *101* (6), 2590–2615.

Kendall, Chad W and Constantin Charles, "Causal narratives," *Working paper*, 2022.

Milgrom, Paul R., "Good News and Bad News: Representation Theorems and Applications," *Bell Journal of Economics*, Autumn 1981, *12* (2), 380–391.

Schwartzstein, Joshua and Adi Sunderam, "Using models to persuade," *American Economic Review*, 2021, *111* (1), 276–323.

Spiegler, Ran, "Bayesian networks and boundedly rational expectations," *The Quarterly Journal of Economics*, 2016, *131* (3), 1243–1290.

_ , "Can agents with causal misperceptions be systematically fooled?," *Journal of the European Economic Association*, 2020, *18* (2), 583–617.