Competing in Tailored Stories

Chiara Aina[*] Marta Mojoli[†]

Extended abstract

Research question. Data rarely speak for themselves: how people react to a piece of evidence depends on the *mental model* they use to interpret it. Politicians, experts, firms, and advocates therefore compete not only over facts but over the stories that explain them. A growing literature studies persuasion through models (Schwartzstein and Sunderam, 2021), in which a sender proposes an interpretation of the data and a naive receiver adopts the one that best fits what she observes. When a single sender can tailor his story to the realized data, he can move beliefs well beyond what Bayesian persuasion allows (Aina, 2026). Most real-world narratives, however, are proposed in competition. In the literature, competition in Bayesian persuasion leads senders to provide the true model (Kamenica and Gentzkow, 2011), and instead competition in ex post model provision leads to neutralize each other's interpretation (Schwartzstein and Sunderam, 2021). This paper asks: *what happens when two senders must propose their models before the data are realized, and the receiver, after observing the signal, adopts the model under which that signal is most likely?* Does ex ante competition discipline manipulation, or can it amplify it?

Setting. A binary state $\omega \in \{\omega_1, \omega_2\}$ with common prior μ_0 generates a binary signal $s \in \{s_1, s_2\}$ through a fixed true model. A *model* m specifies signal distributions conditional on each state; given s, it induces a posterior μ_s^m and has *fit* $\varphi_s^m = \Pr^m(s)$, the likelihood of the observed signal under m. Once observed the signal, the receiver picks the model with the highest explanatory power, that is, the one with the highest fit. This selection rule is supported by experimental evidence showing that, when exposed to competing models of the same data, people tend to adopt the one that best explains what they observed (Aina and Schneider, 2025). There is an equivalence between models and *Bayes-consistent* vectors of posteriors, i.e., vectors $\boldsymbol{\mu} = (\mu_{s_1}, \mu_{s_2})$ that average back to the prior for some weights φ. The set of such vectors sharing the same weights is an *isofit*. Bayes-consistency is a weaker requirement than Bayes plausibility.

[*]Universitat Pompeu Fabra and Barcelona School of Economics. E-mail: chiara.aina@upf.edu.

[†]Department of Economics, University of Zurich. E-mail: marta.mojoli@econ.uzh.ch.

Two senders, a and b, know the true model and each other's preferences. Before the signal is realized, each proposes one model. Upon observing s, the receiver adopts the proposed model with the highest fit given s and updates accordingly. Crucially, selection happens signal by signal, so the receiver may adopt a's model after s_1 and b's after s_2. Ties in fit are broken by each sender's preferred rule with probabilities p and $1 - p$. Sender i evaluates the resulting vector of posteriors under the true signal distribution, $V_i(\boldsymbol{\mu}) = \sum_s \mathrm{Pr}^t(s)\, U_i(\mu_s(\omega_1))$, with U_i continuous, homogeneous of degree one, and strictly monotone. We start by studying pure-strategy equilibria, to then extend to mixed-strategy equilibria, and allowing the senders to send more than one model each.

Mechanism. Because the receiver selects the best-fitting model conditional on each signal realization, the vector of beliefs induced by a *pair* of models need not be Bayes-consistent: it can move beliefs in the same direction after every signal realization. Competition therefore enlarges the set of feasible vectors of posteriors relative to a sender restricted to a single model. The analysis characterizes which of these vectors each sender can attain given the rival's model, and which pairs of models survive mutual best replies.

Results.

1. **Feasibility.** Given a rival model m, we characterize the set of vectors of posteriors a sender can induce with another. It consists of vectors that keep m's posterior on one signal and replace the other with any posterior compatible with m's isofit, together with the isofit of m itself.
2. **Structure of best replies.** Against any non-degenerate model, a best reply (i) must be adopted after at least one signal (proposing a model that is never adopted is never optimal because inconsequential), and (ii) is a *boundary* model, i.e., one that rules out a state after some signal. Against a boundary model, every best reply is *conflicting* (it moves beliefs in the opposite direction on each signal), and the conflicting boundary model with the *same* fit is always a best reply. Against a degenerate model, there is a best reply that is boundary with the same fit. These results reduce the search for equilibria to pairs of boundary, conflicting models.
3. **Aligned senders.** When both senders want to push beliefs toward the same state $\tilde{\omega}$, every non-degenerate equilibrium consists of boundary, conflicting models, and the equilibrium set depends on whether the senders' first best, posterior 1 on $\tilde{\omega}$ after every signal, is attainable:
   - if $\mu_0(\tilde{\omega}) \leq 1/2$, the equilibria are exactly the pairs of boundary, conflicting models with the same fit;
   - if $\mu_0(\tilde{\omega}) > 1/2$, the equilibria additionally include pairs with *different* fit, with $\varphi_{s_1} \in (1 - \mu_0(\tilde{\omega}), \mu_0(\tilde{\omega}))$.

     All of these implement the first best: the receiver becomes certain of $\tilde{\omega}$ regardless of the signal.

   Competition between aligned senders does not discipline persuasion. By proposing conflicting models, each tailored to one signal, the two senders jointly produce Bayes-inconsistent beliefs that no single model could induce alone, and reach the maximal manipulation available to a monopolist who can tailor models to every signal realization.

4. **Misaligned senders.** When senders want to move beliefs in opposite directions, ties are resolved against one of them with positive probability. For any interior tie-breaking probability $p \in (0, 1)$, no pair of non-degenerate, conflicting models with the same fit is an equilibrium because there is always a sender with an incentive to deviate. Across all regions of the fit, one sender profits from a small perturbation of fit that breaks the tie in his favor on the relevant signal. Since best replies to boundary models are conflicting, this points to the non-existence of pure-strategy equilibria under misaligned preferences, reflecting an unraveling race to over-fit the rival's story. This leads to cycles in the models proposed by the senders, and consequentially on the vectors of posteriors that the receiver can hold.

Contribution. The paper extends model persuasion (Schwartzstein and Sunderam, 2021; Aina, 2026) to competition between senders. Unlike Bayesian persuasion with multiple senders, where competition tends to increase information revelation, competition in models can *increase* manipulation when senders' interests are aligned: two conflicting models can accomplish what a single model cannot. When interests are misaligned, the fit-based selection rule generates strong incentives to undercut the rival's fit, and stable pure-strategy outcomes fail to exist. Enlarging the set of models that a sender can provide, competition leads to the opposite effect. Multiple models are used to block this unraveling principle with the consequence that the receiver would stay at the prior.

The failure of stable outcomes under misaligned preferences echoes two recent papers on political competition over narratives, which use different notions of narratives and different techniques, but also their theoretical prediction leads to the emergence of cycles. In Montiel Olea and Prat (2025), candidates announce ideologies, statistical descriptions of a commonly observed data-generating process, and voters weigh how well an ideology fits the data against how simple it is. In Jehiel and Moisson (2026), politicians propose worldviews, partitions of state–policy pairs into analogy classes, and voters with heterogeneous cognitive ability vote for the best-fitting worldview they can understand. In their dynamic model, the complexity of the winning worldview cycles: it rises step by step and then collapses back to the simplest worldview. In both papers, cycles come from a trade-off between fit and simplicity over accumulated data. In ours, the receiver faces no complexity constraint and interprets a single signal realization. Here the instability comes from how the best-fitting model is selected signal by signal, which gives misaligned senders an incentive to keep outbidding their rival's fit.

These papers all together show that because narratives are multidimensional objects, this trade-off in dimensions is exploited by competition in a way that leads to cycles in communication by senders. In our paper the same phenomenon arises, not because of a trade-off between fit and simplicity, or the audience's limited ability to understand complex narratives, but simply because a model must be proposed before the signal is realized. Cycles, in other words, need no cognitive limits: committing to a story before the facts are known is enough.

References

Aina, Chiara (2026) "Tailored stories," Working Paper.

Aina, Chiara and Florian H. Schneider (2025) "Weighting competing models," Working Paper, revise and resubmit at *Econometrica*.

Jehiel, Philippe and P.-H. Moisson (2026) "Political cycles over worldviews: Increments and backlashes in complexity," Working Paper.

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

Montiel Olea, José Luis and Andrea Prat (2025) "Competing ideologies: Fit, simplicity, and fear."

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