MASTERARBEIT | MASTER'S THESIS

Titel | Title

When many speak, less is said: how costly signals reduce information in multi-sender Bayesian Persuasion

verfasst von | submitted by

Lars Thieme B.Sc. MA

angestrebter akademischer Grad | in partial fulfilment of the requirements for the degree of

Master of Science (MSc)

Wien | Vienna, 2025

Studienkennzahl lt. Studienblatt | Degree programme code as it appears on the student record sheet:	UA 066 953
Studienrichtung lt. Studienblatt | Degree programme as it appears on the student record sheet:	Masterstudium Research in Economics and Finance
Betreut von | Supervisor:	Univ.-Prof. Dipl.-Math. Karl Schlag PhD

Acknowledgments

This thesis marks the end of a long journey of the challenging yet immensely fulfilling time I had as a student at the University of Vienna. I could never have gotten this far without the invaluable support I constantly received from everyone around me.

I want to thank my Supervisor, Karl Schlag, for everything he taught me and for believing in me when I doubted myself .

I am grateful to my friends for uplifting me in darker moments and making the lighter times shine even brighter.

Above all, I want to thank my parents. They have always been the first to cheer me on, persistently encouraging me in all of my pursuits.

Abstract

Do decision-makers always learn more when receiving information from multiple sources? Kamenica and Gentzkow (2017a) show that in a Bayesian persuasion framework, increasing the number of senders increases the informativeness in equilibrium. We challenge this result by introducing costly signals into a multi-sender persuasion model and show how two competing senders provide less information than a single sender. Costly competition introduces a new trade-off, making pure-strategy equilibria infeasible. We characterize a mixed-strategy ε-equilibrium and show how competition reduces the incentives to reveal information. Our findings supply reason for caution when regulators aim to extract more information from competing sources.

Kurzfassung

Lernt eine Entscheidungsträgerin grundsätzlich mehr, wenn sie Informationen aus mehreren Quellen erhält? Kamenica and Gentzkow (2017a) zeigen, dass in einem "Bayesian Persuasion" Modell die Erhöhung der Anzahl der Sender die Informativität im Gleichgewicht erhöht. Wir stellen dieses Ergebnis in Frage, indem wir kostspielige Signale einführen und zeigen, dass zwei konkurrierende Sender weniger Informationen liefern als ein einzelner Sender. Der kostspielige Wettbewerb führt einen neuen Zielkonflikt ein, der Gleichgewichte mit reinen Strategien unmöglich macht. Wir charakterisieren ein ε-Gleichgewicht mit gemischten Strategien und zeigen, wie der Wettbewerb die Anreize zur Offenlegung von Informationen verringert. Unsere Ergebnisse geben Anlass zur Vorsicht, wenn die Massnahmen darauf abzielen, mehr Informationen aus konkurrierenden Informationsquellen zu gewinnen.

1. Introduction

Does a decision-maker always learn more when receiving information from multiple senders compared to receiving it from just one? Conventional wisdom suggests yes, as we are told to "consult multiple sources", to "get a second opinion", or to "hear both sides". Observing information from different senders is especially valuable if the individual sources of information are biased or selective. Examples of selective information provision include situations where information is provided by interested parties who aim to persuade the decision-maker to choose a specific course of action.

Consider a criminal trial, in which a prosecutor and a defense attorney present information to persuade the judge to convict or acquit. Intuitively, we would expect the presence of both senders to be beneficial for the judge. For example, we might expect that the information coming from the defense attorney prevents the prosecutor from misrepresenting the case. Kamenica and Gentzkow (2017a) have formalized this line of thought, showing that increasing the number of senders in a Bayesian persuasion framework can only increase the informativeness of the equilibrium. In particular, they argue that in the specific setting of a judge, who receives information from a prosecutor and a defender, the unique equilibrium outcome is full revelation by both senders.

The objective of this thesis is to challenge the above claim that more senders always result in more information in Bayesian persuasion models. We characterize a setting where two senders (Prosecutor and Defender) jointly reveal less information than a single sender (Prosecutor) would in isolation. While it remains true that the additional sender cannot directly reduce the informativeness of Prosecutor's signal, the mere presence of a competitor alters Prosecutor's incentives. In particular, this can significantly reduce the informativeness of Prosecutor's optimal signal.

In our model, we introduce the assumption that sending a more informative signal incurs increased costs for the senders. We motivate the inclusion of costs in two ways. Firstly, we find it odd that competition should force *both* senders to fully reveal, as in the literature benchmark. To that extent, we use costs as instruments to incentivize each sender to provide less information if the other sender already provides a lot. Furthermore, we argue that costs are also a reasonable assumption in many persuasion contexts.

Information provision in persuasion settings works by choosing signals, which are commonly interpreted as experiments or mechanisms that generate information. In our criminal trial example, a signal could be a selection of conducted investigations consisting of various interviews and forensic tests. In the case of lobbying, a signal could be a commissioned study, and in the case of advertising, it could be a series of lab experiments that test product quality. The bottom line is that all of these procedures are associated with costs and are plausibly more costly the more revealing they are.

The main consequence of the additional trade-off between a more informative signal and higher costs is that the two-sender persuasion model with costs has no equilibrium in pure strategies. We provide a sketch of a proof and characterize an ε–equilibrium in mixed strategies. We arrive at this ε–equilibrium by discretizing the strategy space and applying a Lemke-Howson algorithm (Lemke and Howson (1964)) to find a Nash equilibrium in discrete space and refine it in continuous space. As the measure for the amount of information revealed, we use the expected reduction in entropy resulting from the joint signals of both senders. When comparing the results obtained in one-sender models against two-sender models with and without costs, we find that in the absence of costs, including a second sender drastically increases informativeness. In contrast, on our model with costly signals, informativeness *decreases* when a second sender is added.

Our findings suggest that we should be more cautious when expecting additional information from considering multiple competing sources. This is particularly important in the context of growing public concern about misinformation and the reliability of information sources. By clarifying the conditions under which more senders lead to less informative communication, our work offers insight into how specific information environments may hinder rather than help public understanding, an issue central to democratic discourse, scientific communication, and institutional trust.

We emphasize that this thesis should not be viewed as a self-contained exercise in theoretical reasoning, as the motivation is to contribute to our understanding of strategic information provision in the real world.

1.1. Literature

The two most relevant strands of the literature, to which we connect, are persuasion models involving multiple senders and persuasion models involving costly signals.

Bayesian persuasion coined by Kamenica and Gentzkow (2011) has received considerable attention with various applications and extensions, which Bergemann and Morris (2016) and Kamenica (2019) have surveyed. Surprisingly, a direction that has received comparatively little attention until recently are persuasion models that feature *multiple* senders. This is odd, because in the popular applications used in the persuasion literature, such as criminal trials, lobbying, or advertising, we would expect to find multiple senders.

The benchmark models of multi-sender persuasion are Kamenica and Gentzkow (2017a,b). Among their results is that increasing the number of senders weakly increases the informativeness of the equilibrium. The three conditions they need to arrive at this result are that signals can be arbitrarily correlated (1), senders move simultaneously (2), and play pure strategies (3). Li and Norman (2018) show that these three conditions are indeed both necessary and sufficient, as they provide three short counterexamples, where they separately relax each one of those conditions.

Moreover, other contributions have previously discussed settings where multiple senders decrease informativeness. Au and Kawai (2020) do so by constructing a model in which each sender can only reveal distinct parts of the state space. Li and Norman (2021) as well as Koessler et al. (2022) reach the conclusion of less information from multiple senders in the context of sequential moves persuasion models.

The bottom line is that we are not the first to provide an answer to the challenge of whether more senders always provide more information. However, our contribution is to show that we can arrive at this conclusion without explicitly relaxing one of the three conditions that Kamenica and Gentzkow specified. In our model, we assume conditional independence of signals (which is nested in the arbitrarily correlated condition), senders are allowed to play pure strategies, they move simultaneously, and each sender can fully reveal. Yet, we still find that two senders provide less information than a single one.

There are further papers that argue in favor of *increased* informativeness under multiple senders, but require that the receiver possesses additional power to achieve these results. Gradwohl et al. (2022) and Boleslavsky and Cotton (2018) present models in which the receiver can select whether to receive information from multiple senders or just one. Wu (2023) consider a sequential multi-sender persuasion model, in which they find a fully revealing equilibrium. When testing single against multiple-sender setups experimentally, Wu and Ye (2023), find that multiple senders do reveal more information, but not quite full revelation.

The idea of including costly signals in persuasion models is not new. Our implementation of adding a simple cost function, which increases in the amount of information revealed by a sender, to the payoff of each sender is adopted from Kamenica and Gentzkow (2014). What is new is that we extend this framework to multiple senders.

Several single-sender papers embed cost structures that create richer strategic incentives than merely linking more information to higher costs. Degan and Li (2021) separate costly information acquisition from signal transmission, so that the receiver observes both the precision chosen and the message sent, which allows costs to serve as a credibility signal. Nguyen and Tan (2021) introduce signal-dependent message costs that can make it more costly to misrepresent the state. How these more nuanced cost mechanisms would play out with multiple competing senders remains an open question.

The remainder of the thesis is structured as follows. In chapter 2, we begin with a description of the model containing the definition of all relevant terms, as well as a brief illustrative example. In chapter 3, we solve the model by characterizing an ε−equilibrium, before we analyze the properties of the ε−equilibrium concerning the amount of revealed information and present the results in chapter 4. We provide a discussion of our results, including policy implications and limitations of our model section 5.1. Lastly, we end with a conclusion in chapter 6.

2. Model

2.1. Setup

For ease of exposition, we present the model environment in terms of a courtroom setting, consisting of a Prosecutor, a Defender, and a Judge. This means that our work is a direct extension of the seminal example of the original persuasion model by Kamenica and Gentzkow (2011). Moreover, we argue that the courtroom is a particularly well-suited example for our purpose, as it closely resembles our setting of two competing senders who each incur costs when they provide more information. In the following, we describe the model environment and define all relevant terms.

Let $\omega \in \{I, G\}$ be a binary unknown state denoting whether the accused is innocent or guilty. There is a single receiver (Judge) and two senders indexed by $i \in \{P, D\}$ denoting Prosecutor and Defender, respectively, who all share a common prior $\mu_0 = \Pr(\omega = G) = 0.3$. For easier comparison and interpretation, we choose the prior to be the same as in Kamenica and Gentzkow (2011).

Following the literature, we model the Judge purely mechanically, in that he decides based on a predetermined cutoff rule. Judge convicts if and only if the observed posterior resulting from the senders' signal realizations is at least 0.5 and acquits otherwise. This is equivalent to saying the Judge's preferences are to match his action to the state. Since our $\varepsilon-$equilibrium will be in mixed-strategies over a continuous support, a posterior realization exactly on the threshold has probability zero, which means that it does not matter if we use a strong or weak inequality in the definition of the decision rule.

Senders choose a binary signal that consists of distributions over signal realizations $m_i \in \{g_i, i_i\}$, which means that a realization conveys either a "guilty" or an "innocent" message. All players observe both the chosen signal and the resulting realization.

The timing of the model is straightforward. Both senders simultaneously choose a signal. Judge observes both signals and signal realizations, calculates the resulting posterior, and decides according to the cutoff rule. Since senders move simultaneously and the judge follows a fixed cutoff rule, we solve for a Nash equilibrium over the senders' strategy profiles.

2.2. Interim beliefs

The concept of interim beliefs is adopted from Kamenica and Gentzkow (2017a). In a single-sender model, each signal realization yields a unique posterior, implying that the sender directly chooses a distribution of posteriors, subject to feasibility constraints. In the two-sender model, no single sender can determine the posterior alone, as it depends on the signal realizations of both senders. Instead, we model the senders as choosing "interim beliefs", which are the hypothetical posteriors that would be induced if Judge only observed one sender's signal. Consequently, sender i induces a distribution over a pair of interim beliefs $\mu_i = (\mu_i^L, \mu_i^H)$ with $\mu_i^L \in [0, \mu_0], \mu_i^H \in [\mu_0, 1]$, which are the hypothetical posteriors for observing only a "guilty" signal realization or "innocent" realization from sender i alone. The interim beliefs are denoted with superscripts "H" and "L", since the Bayes-plausibility constraint implies that one interim belief has to be (weakly) higher than prior and one (weakly) lower than the prior:

$$\mu_i^H = \Pr(\omega = G|g_i) = \frac{\mu_0 \Pr(g_i|G)}{\mu_0 \Pr(g_i|G) + (1-\mu_0)\Pr(g_i|I)}$$
$$\mu_i^L = \Pr(\omega = G|i_i) = \frac{\mu_0 \Pr(i_i|G)}{\mu_0 \Pr(i_i|G) + (1-\mu_0)\Pr(i_i|I)}.$$

Using $\Pr(g_i|G) = 1-\Pr(i_i|G)$ and $\Pr(g_i|I) = 1-\Pr(i_i|I)$, we can recover the corresponding signals in terms of the interim beliefs as:

$$\Pr(g_i \mid G) \; = \; \frac{\mu_i^H\left(\mu_0 - \mu_i^L\right)}{\mu_0\left(\mu_i^H - \mu_i^L\right)}, \qquad \Pr(g_i \mid I) \; = \; \frac{\left(\mu_0 - \mu_i^L\right)\left(1 - \mu_i^H\right)}{\left(1-\mu_0\right)\left(\mu_i^H - \mu_i^L\right)}.$$

Each sender's choice of interim beliefs has to fulfill the Bayes-plausibility constraint, which requires that the distribution over interim beliefs is, in expectation, equal to the prior. In the binary setting, we express the constraint as:

$$\mathbb{E}[\mu_i] = \mu_0 \;\Leftrightarrow\; \lambda_i\mu_i^H + (1-\lambda_i)\mu_i^L, \text{ for } \lambda_i \in [0,1].$$

λ_i and $(1-\lambda_i)$ are the frequencies with which the two interim beliefs are induced. These frequencies directly follow from the chosen interim beliefs:

$$\lambda_i\mu_i^H + (1-\lambda_i)\mu_i^L = \mu_0 \Leftrightarrow \lambda_i = \frac{\mu_0 - \mu_i^L}{\mu_i^H - \mu_i^L}.$$

Moreover, we may also think of λ_i as the frequency of player i's guilty signal realization,

$$\lambda_i = \Pr(g_i) = \Pr(g_i|G)\Pr(G) + \Pr(g_i|I)\Pr(I).$$

2.3. Costs

The driving mechanism in our model is the addition of costly signals in the multi-sender persuasion framework. We assume that more revealing signals are more costly, as in Kamenica and Gentzkow (2014). We choose signal costs as the squared differences between the prior and the interim beliefs, weighted by the frequency λ_i with which these interim beliefs are induced and an exogenous scaling parameter k:

$$c(\mu_i^H, \mu_i^L) = k\left[\lambda_i(\mu_i^H - \mu_0)^2 + (1-\lambda_i)(\mu_i^L - \mu_0)^2\right].$$

Note that plugging in the expression for λ_i in terms of the interim beliefs reveals that our cost function is equivalent to the product of the distances between interim beliefs and the prior:

$$c(\mu_i^H, \mu_i^L) = k(\mu_0 - \mu_i^L)(\mu_i^H - \mu_0).$$

Our setup allows an uninformative and therefore costless signal. In particular, costs are zero if at least one interim belief equals the prior. If either interim belief equals the prior, then $\lambda_i \in \{0, 1\}$ and the signal becomes completely uninformative and therefore costless.

Before we move on, we briefly discuss the scaling parameter k. To refute the result that two senders uniquely reach full revelation by both senders in equilibrium, an arbitrarily small value of k suffices. The reason is that the unique best response to full revelation by the other sender is to play the uninformative signal to incur zero costs, which means that there is no equilibrium at mutual full revelation. Moreover, full revelation is generally not a best response to the uninformative signal. We pick up on this argument when we present our results in section 4.2. In the remainder of the thesis, we consider a setting with a moderate cost parameter of $k = 10$.

2.4. Posteriors

We assume that Prosecutor's and Defender's signals are conditionally independent given the state, which delivers closed-form expressions for the posteriors. Moreover, this assumption is nested within the assumption of arbitrarily correlated signals, as presented in our literature benchmark of Kamenica and Gentzkow (2017a). We point this out explicitly to argue that our results are comparable. A caveat is that real-world attorneys also draw on shared evidence, for example, a forensic report, which can induce correlated signals and thus violate our assumption. We abstract away from this possibility to keep the model tractable.

In our model, the general expression of a posterior resulting from two signal realizations $m_i \in \{g_i, i_i\}$ is given by:

$$\Pr(\omega = G|m_P, m_D) = \frac{\mu_0 \cdot \Pr(m_P|G) \cdot \Pr(m_D|G)}{\mu_0 \cdot \Pr(m_P|G) \cdot \Pr(m_D|G) + (1-\mu_0) \cdot \Pr(m_P|I) \cdot \Pr(m_D|I)}.$$

Since we have two senders with a binary signal each, there are four distinct posterior realizations, which result from the 2×2 combinations of signal realizations. We express these posterior realizations in terms of the senders' chosen interim beliefs and their frequencies f_{PD} as:

$$\begin{aligned}
\hat{\mu}(\mu_P^H, \mu_D^H) &= \frac{(1-\mu_0)\mu_P^H \mu_D^H}{\mu_P^H \mu_D^H - \mu_P^H \mu_0 - \mu_D^H \mu_0 + \mu_0}, \text{ with } f_{HH} = \lambda_P \cdot \lambda_D + \mathrm{Cov}(\cdot) \\
\hat{\mu}(\mu_P^H, \mu_D^L) &= \frac{(1-\mu_0)\mu_P^H \mu_D^L}{\mu_P^H \mu_D^L - \mu_P^H \mu_0 - \mu_D^L \mu_0 + \mu_0}, \text{ with: } f_{HL} = \lambda_P \cdot (1-\lambda_D) - \mathrm{Cov}(\cdot) \\
\hat{\mu}(\mu_P^L, \mu_D^H) &= \frac{(1-\mu_0)\mu_P^L \mu_D^H}{\mu_P^L \mu_D^H - \mu_P^L \mu_0 - \mu_D^H \mu_0 + \mu_0}, \text{ with: } f_{LH} = (1-\lambda_P) \cdot \lambda_D - \mathrm{Cov}(\cdot) \\
\hat{\mu}(\mu_P^L, \mu_D^L) &= \frac{(1-\mu_0)\mu_P^L \mu_D^L}{\mu_P^L \mu_D^L - \mu_P^L \mu_0 - \mu_D^L \mu_0 + \mu_0}, \text{ with: } f_{LL} = (1-\lambda_P) \cdot (1-\lambda_D) + \mathrm{Cov}(\cdot),
\end{aligned}$$

where,

$$\mathrm{Cov}(\mu_P^L, \mu_P^H, \mu_D^L, \mu_D^H, \mu_0) = \frac{\lambda_P \lambda_D (\mu_P^H - \mu_0)(\mu_D^H - \mu_0)}{\mu_0(1-\mu_0)}.$$

For interpretation of posteriors in terms of interim beliefs, note that they correspond to posteriors resulting from combinations of signal realizations of the senders' signals:

$$\begin{aligned}
\hat{\mu}(\mu_P^H, \mu_D^H) &= \Pr(\omega = G|g_P, g_D) \\
\hat{\mu}(\mu_P^H, \mu_D^L) &= \Pr(\omega = G|g_P, i_D) \\
\hat{\mu}(\mu_P^L, \mu_D^H) &= \Pr(\omega = G|i_P, g_D) \\
\hat{\mu}(\mu_P^L, \mu_D^L) &= \Pr(\omega = G|i_P, i_D).
\end{aligned}$$

2.5. Payoffs

In addition to incurring the costs determined by the choice of interim beliefs, Prosecutor gets a payoff of one if Judge convicts, i.e., posterior $\hat{\mu} \geq 0.5$, and a payoff of zero minus costs else. Analogously, Defender has a preference for acquittal, which is given by $\hat{\mu} < 0.5$. The senders' gains from Judge's action are zero-sum; however, since the senders are free to choose how revealing their signals are, they incur different costs, which means that the game overall is not zero-sum.

Since the posterior is a random variable with four possible realizations, the payoff of a sender is given by the expected probability of winning the case, which we express as the sum of frequencies of posterior realizations larger than or equal to the 0.5 threshold, minus the cost of information provision:

$$U_P = \left[\sum_{\hat{\mu}_{PD} \geq 0.5} f_{PD} \right] - c(\mu_P^H, \mu_P^L) \qquad U_D = \left[\sum_{\hat{\mu}_{PD} < 0.5} f_{PD} \right] - c(\mu_D^H, \mu_D^L) \qquad P, D \in [L, H].$$

3. Solving the model

3.1. Single-sender benchmarks

Before we solve for our ε–equilibrium in the two-sender model, we briefly present the single-sender equilibria. On the one hand, they serve as useful benchmarks for comparing our multi-sender model outcomes, and on the other hand, they provide a concise illustration of the game's mechanisms.

Only Defender A persuasion model that only includes Defender is straightforward. Given that the prior is below the conviction threshold, the default action of the judge is to acquit in the absence of any additional information. Because this is already the preferred outcome for Defender, the optimal strategy is to provide no information, thereby incurring zero costs and guaranteeing a payoff of 1.

Only Prosecutor To solve for the optimal signal of Prosecutor in a single-sender model, we draw on the arguments by Kamenica and Gentzkow (2014), given that the only difference in our model is a slightly different functional form for the cost function.

Prosecutor will choose such that the high posterior is at least equal to 0.5, to ensure conviction whenever a "guilty" signal realization occurs. Expected payoff is therefore given as:

$$U_P(\mu_P^H \geq 0.5,\ \mu_P^L) = \lambda_P - c(\mu_P^H, \mu_P^L).$$

Moreover, we know that the constraint has to be binding because U_P is strictly decreasing in μ_P^H for $\mu_P^H > 0.5$. An intuitive explanation is that Prosecutor does not gain any additional benefit if Judge's posterior is any larger than the threshold, because the decision is determined by the cut-off only.

Prosecutor's simplified problem is therefore given as:

$$\max_{\mu_P^L} U_P(\mu_P^L, \mu_P^H = 0.5) = \frac{\mu_0 - \mu_P^L}{0.5 - \mu_P^L} - k(0.5 - \mu_0)(0.3 - \mu_P^L)),$$

which yields the solution:

$$\mu_P^{L^*} = 0.5 - \frac{1}{\sqrt{k}}.$$

Note that this solution is only in the interior for specific values of the parameters (μ_0, k). For example, if k is too large, persuasion is no longer cost-efficient for Prosecutor, meaning that the optimal signal is to provide no information at all.

For an interior solution, we need to ensure that the expected payoff is positive:

$$U_P^* = \frac{\mu_0 - \mu_P^L}{0.5 - \mu_P^L} - k(0.5 - \mu_0)(0.3 - \mu_P^L) \geq 0.$$

Using $\mu_0 = 0.3$ we obtain the following bounds on $k \in [4, 25]$, which simultaneously ensure $\mu_P^L \in [0, 0.3]$ and $U_P^* \geq 0$. This result also provides some context for our parameter choice, $k = 10$, in the two-sender model. We compare and interpret the equilibrium outcome of this single-prosecutor model with costs to the setting without costs, as well as to the two-sender model, when we present our main results in section 4.1.

3.2. No equilibrium in pure strategies

We begin with a sketch of a proof that there is no equilibrium in pure strategies. Intuitively, a sender can make an infinitesimal tweak to his signal to push a posterior lying just on one side of the 0.5 cutoff to the other side, while incurring virtually no additional cost. Accordingly, the proof consists of two steps. First, we show that Prosecutor's best response is always such that at least one of the posteriors is precisely at the 0.5 threshold. The second step is to show that whenever one of the posteriors equals the threshold, Defender has a profitable pure strategy deviation, namely, nudging the posterior below the threshold.

Step 1: Prosecutor optimality requires binding threshold

Suppose $\hat{\mu}_{HL} > 0.5$. Then there exists $\tilde{\mu}_P^H < \mu_P^H$ such that $\hat{\mu}(\tilde{\mu}_P^H, \mu_D^L) = 0.5$. Since

$$c(\tilde{\mu}_P^H, \mu_P^L) < c(\mu_P^H, \mu_P^L), \quad \lambda_P(\tilde{\mu}_P^H, \mu_P^L) > \lambda_P(\mu_P^H, \mu_P^L) \Leftrightarrow \tilde{f}_{HL} > f_{HL}$$

it follows that:

$$U_P(\tilde{\mu}_P^H, \mu_P^L) > U_P(\mu_P^H, \mu_P^L).$$

In words, if the posterior is strictly above the threshold, it is profitable for Prosecutor to deviate to a lower high interim belief, because this reduces costs and increases the frequency of this posterior realization, both of which leads to a higher payoff. The same argument can be made for situations in which $\hat{\mu}_{HL} < 0.5$, where the best response is then to choose such that $\hat{\mu}_{HH} = 0.5$.

Step 2: profitable Defender deviation exists

Assume a pure strategy equilibrium (μ_P^*, μ_D^*) exists. Then:

$$\exists\, x, y \in \{H, L\} \text{ such that } \hat{\mu}_{xy} = 0.5$$

Defender can then deviate marginally (e.g. $\mu_D^L \to \mu_D^L - \varepsilon$) to shift $\hat{\mu}_{xy} < 0.5$, yielding:

$$\Delta U_D > 0, \quad \Delta c_D \to 0$$

Thus, (μ_D^*, μ_P^*) is not a mutual best response.

Lastly, there is also the possibility that Prosecutor best responds by providing no information. However, this cannot be an equilibrium either, because the unique Defender best response to no information by Prosecutor is to provide no information as well. However, in that case, Prosecutor has an incentive to deviate because if Defender provides no information at all, the game then becomes identical to the single-Prosecutor game, which we have already solved.

3.3. Mixed strategies

Following the result that there is no Nash equilibrium in pure strategies, we search for an equilibrium in mixed strategies. A mixed strategy equilibrium in our game is defined in the standard way. For each sender $i \in \{P, D\}$, recall that the strategy space S_i is given by a subset of $\mathbb{R}^2$, which represents pairs of induced interim beliefs:

$$S_i = \left\{(\mu_i^L, \mu_i^H) \in (0,1)^2 \ : \ \mu_L^i < \mu_0 < \mu_H^i\right\}.$$

A mixed strategy σ_i for player i is a probability measure over S_i. For a pure strategy $s_i \in S_i$ and a mixed strategy σ_{-i} of the opponent we define

$$U_i\Big(s_i; \sigma_{-i}\Big) = \int_{S_{-i}} U_i(s_i, s_{-i})\, d\sigma_{-i}(s_{-i}),$$

and player i's equilibrium pay-off as

$$U_i^*(\boldsymbol{\sigma}) = \iint_{S_P \times S_D} U_i(s_P, s_D)\, d\sigma_P(s_P)\, d\sigma_D(s_D).$$

A strategy profile $\boldsymbol{\sigma}^* = (\sigma_P^*, \sigma_D^*)$ is a mixed-strategy Nash equilibrium if, for both senders it holds that:

$$U_i(\sigma_i^*, \sigma_{-i}^*) \geq U_i(\sigma_i, \sigma_{-i}^*) \ \forall \ \sigma_i.$$

Given that the strategy space is infinite on a bounded two-dimensional plane and the payoff functions have discontinuities, the exact location of which depends on the strategies of both players, we do not solve for an equilibrium analytically. Instead, we proceed to approximate an ε–equilibrium of the game, for which we start with the following simplification. We discretize the continuous strategy space by allowing interim beliefs to move only in steps of 0.02, meaning that the strategy space consists of $[\frac{0.3}{0.02}] \times [\frac{0.7}{0.02}] = 525$ possible pure strategies for each sender. This increases tractability significantly.

Moreover, in the discrete game, we can numerically apply iterated elimination of strictly dominated strategies by computing the payoff for each of the 525×525 cells in the normal form game. After six iterations, the algorithm converges and yields 66 pure strategies for Prosecutor and 47 for Defender. We illustrate the discrete strategy space, as well as the reduced strategy space after iteratively removing strictly dominated strategies below in Figure 3.1.

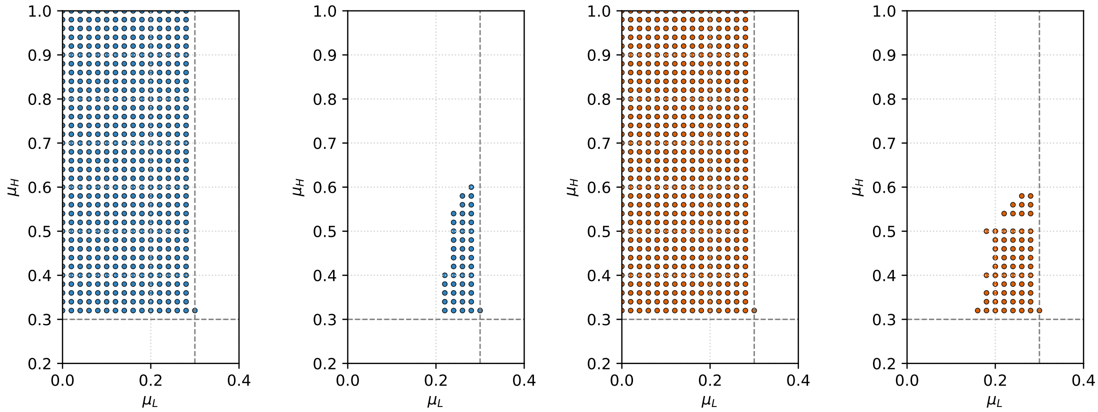

(a) Defender full (b) Defender reduced (c) Prosecutor full (d) Prosecutor reduced

Figure 3.1.: Comparison of the full vs. reduced strategy spaces at grid size 0.02

We solve this discretized normal form game with the reduced strategy set by applying a Lemke-Howson algorithm in Gambit (Python) (Lemke and Howson (1964); Savani and Turocy (2025)). Note that we cannot further reduce the grid size as the Lemke-Howson algorithm does not converge if the payoff differences between pairs of pure strategies of a sender become too small. More robust algorithms, such as support enumeration, could theoretically handle these difficulties of smaller grids, however, support enumeration algorithms are highly sensitive to the number of pure strategies, which makes them unsuitable in our case. We illustrate the equilibrium mixed strategy profile in Table 3.1.

σ_P			σ_D		
μ_P^L	μ_P^H	$\Pr\big((\mu_P^L, \mu_P^H)\big)$	μ_D^L	μ_D^H	$\Pr\big((\mu_D^L, \mu_D^H)\big)$
0.26	0.46	0.0394	0.24	0.32	0.1062
0.26	0.48	0.0887	0.24	0.34	0.2644
0.26	0.50	0.1492	0.26	0.34	0.2384
0.26	0.54	0.2411	0.28	0.34	0.2403
0.26	0.56	0.3375	0.28	0.36	0.0617
0.26	0.58	0.1441	0.30	0.32	0.0889

Table 3.1.: Discrete space equilibrium: supports and probability weights

In the discretized game, this equilibrium results in payoffs:

$$U_P^* \approx 0.0130 \qquad\qquad U_D^* \approx 0.8777$$

These six discrete support elements with respective probability weights for each sender constitute a first approximation of an equilibrium in continuous space. To get a measure of how close the discrete strategies are to an equilibrium in continuous space, we simulate the mixed strategies against a continuum of possible pure strategy deviations in the whole continuous strategy space. In particular, we draw $10,000$ pure-strategy pairs (μ_i^L, μ_i^H) for each sender and compute the expected payoff against the opponent's mixed strategy. For each sender, we denote the difference between the highest pure strategy payoff $\tilde{U}_i$ against the equilibrium payoff U_i^* by $\varepsilon_i = \max\{\tilde{U}_i - U_i^*\}$. For the reported mix of six discrete support strategies, we find the following deviation magnitudes when paired with pure strategy deviations from the continuous space:

$$\varepsilon_P = 0.0154 \qquad\qquad \varepsilon_D = 0.0223.$$

3.4. ε–Equilibrium in continuous space

The mixed strategy profiles we derived in the previous section are only a crude approximation of an equilibrium in continuous space, because senders are artificially forced into the discrete grid. In contrast, the true equilibrium supports may lie in the space between the discrete grid points.

We use the support and probability weights found in the discretized game equilibrium to construct a closer approximation of an ε-equilibrium in the continuous strategy space, meaning that the goal is to find a pair of mixed strategies that yields lower values for ε_i.

We conjecture that such an ε–equilibrium exists in the continuous neighborhood around the support of the equilibrium in the discretized game. Our approach is therefore to construct the support in continuous space as a small area around each discrete support point. Next, we assign each constructed area the probability weight of the respective support element from the discrete profile. Lastly, we spread the density of each discrete support point uniformly across the constructed neighborhood area, thereby obtaining a fully continuous mixed strategy profile.

Figure 3.2 illustrates the support for both senders in the discretized setting in the left panel and the resulting continuous areas that we constructed around each discrete point in the right panel. Given the shape of the support in the discrete setting, we construct the continuous supports not as uniform areas around each grid point (e.g., circles or squares), but as thin strips. We conjecture these shapes from the location of the discrete support points.

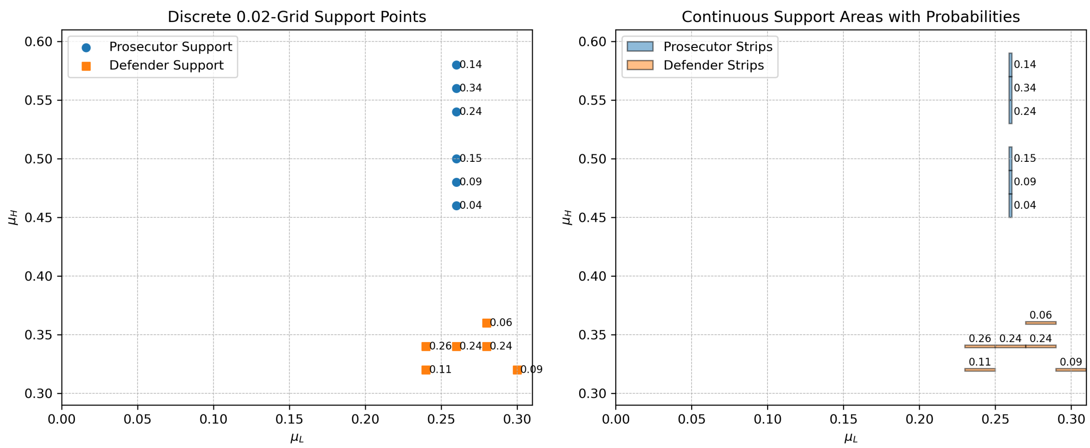

Figure 3.2.: Discrete support (left) and continuous support (right)

Using these continuous support areas, we obtain a closer approximation, as indicated by lower pure strategy deviation magnitudes:

$$U_P^* = 0.0099 \qquad U_D^* = 0.8807$$
$$\varepsilon_P = 0.0061 \qquad \varepsilon_D = 0.0110.$$

These most profitable deviations are obtained from the pure strategy $\tilde{\mu}_P^L = 0.2583, \tilde{\mu}_P^H = 0.4555$ in the case of Prosecutor, and $\tilde{\mu}_P^L = 0.2583, \tilde{\mu}_P^H = 0.4555$, and $\tilde{\mu}_D^L = 0.2919, \tilde{\mu}_D^H = 0.3658$ for Defender. For both senders, these pure strategies lie just outside the support of the discrete strategy profiles. Furthermore, they are close enough that they are included in the mixed strategy profile through the expansion of the discrete points.

This suggests that we are close to an equilibrium, but the support and probability weights are not entirely correct. While the deviations are minor in absolute terms, Prosecutor can increase payoff by roughly 62 percent, which is a sizable deviation. This underlines that the strategy profile is only an approximation of an equilibrium.

For the remainder of the thesis, we treat this mixed-strategy profile as the benchmark outcome of the two-sender persuasion model with costly signals. The next step is to analyze it in terms of the amount of information revealed and draw comparisons between the multi-sender and single-sender model specifications.

4. Results: when two senders say less

4.1. Simulated ε−Equilibrium properties

Given our guess for an ε−equilibrium, we proceed to analyze the outcome of the strategy profile concerning the amount of information that the two senders jointly reveal. To do this, we run the following Monte-Carlo simulation with $N = 10,000$ draws:

1) draw unknown state $\omega \in \{I, G\}$ according to μ_0

2) draw a pure strategy for each sender according to their continuous mixed strategies

3) convert interim beliefs (μ_i^L, μ_i^H) into signals $\Pr(g_i|G)$ and $\Pr(g_i|I)$

4) draw a signal realization for each sender

5) compute the resulting posterior based on the two signal realizations

6) calculate the reduction in entropy between this posterior and the prior

Our measure of signal informativeness is the average reduction in Shannon Entropy $\mathbb{E}\Delta H(x)$. In each simulation round n we calculate the entropy of the realized posterior and compare its entropy to the prior, i.e.

$$\Delta H_n = H(\mu_0) - H(\hat{\mu}_n).$$

Given that the unknown state is binary, the Shannon Entropy measured in Nats is given as:

$$H(x) = -[x \ln(x) + (1-x)\ln(1-x)].$$

Using the outcomes of the simulation of the mixed-strategy profile, we summarize our main results in Table 4.1. In particular, we compare the original single-prosecutor persuasion model Kamenica and Gentzkow (2011) with the full revelation two-sender result from Kamenica and Gentzkow (2017a), as well as the results we derive in this thesis for both the single-sender and two-sender settings. We acknowledge that the comparability of results is limited since the mixed strategy profile we used in the two-sender model with costly signal is not an equilibrium, whereas the remaining three columns represent equilibrium outcomes of the respective model specifications. This is also reflected in the fact that the values in the first three columns refer to analytical expected values, whereas the fourth column displays mean values derived from the Monte-Carlo simulation. For the derivation of the analytical results, we refer to Appendix A.1.3.

	No cost (literature)		**With cost (this thesis)**	
	Single prosecutor	Two senders	Single prosecutor	Two senders
Judge accuracy	0.70	1.00	0.70	0.71
Entropy reduction	0.1950	0.6109	0.0540	0.0240
Prosecutor payoff	0.6	0.3	0.1351	0.0099
Defender payoff	-	0.7	-	0.8807

Table 4.1.: Comparison of outcomes across model specifications

The two dimensions along which the different model specifications can be compared are the number of senders and whether signals are costly. The comparison that matters most for this thesis is the difference in entropy reduction between the two-sender and one-sender models, respectively.

The central result of this thesis is that *under costly signals, moving from a single prosecutor to two competing senders reduces the expected amount of information revealed*. This is the exact opposite effect to the specifications without costs, where two competing senders uniquely fully reveal in equilibrium. We find that the introduction of competition under costly signals reduces the joint informativeness, because it incentivizes senders to choose less informative signals than they would otherwise.

We also observe several other differences between the models, that are interesting to discuss. For a given number of senders, the inclusion of costly signals results in less information being revealed on average. In the case of a single sender, this aligns with the conclusion drawn by Kamenica and Gentzkow (2014). The mechanism is that under costly signals, the sender's payoff function over the posterior becomes increasingly concave, since more extreme posteriors are associated with higher induced costs. A single sender, therefore, faces a trade-off between revealing more information to favor their desired Judge's action and incurring higher costs. Moreover, we also see that comparing the two-sender models, the setting with costly signals is associated with less information. In summary, the reduced information through costly signals in the single-sender model is an application of the existing literature. The outcome in the two-sender model can be seen as an extension of this result.

Examining the sender payoffs, we observe that the single Prosecutor's payoff decreases with the introduction of the cost function and is even worse off in the setting of costly competition. The reason behind the latter is the advantageous position of Defender resulting from the biased prior. Given the prior $\mu_0 = 0.3$, Prosecutor has to induce stronger belief changes to harvest a positive payoff from Judge's decision, which comes at increased costs. Defender is able to exploit the strategically advantageous position, which is reflected in a much higher payoff.

Interestingly, Defender is even better off under costly competition than without costs. This result is again due to the advantageous position given by the Judge's default action. Since there is only little information revealed in this model specification, Judge's decisions will be closer to the prior on average, which benefits Defender.

We also compare the model specifications concerning Judge accuracy, which denotes the share of correct decisions made by the Judge and would be the Judge's payoff if we modeled him as an additional player in this game. Naturally, the Judge is best off in the two-sender model without costs, as under full revelation he makes perfectly accurate decisions. What Kamenica and Gentzkow (2014) have already shown is the Judge is not worse off under costly signals in these settings, because the costs only restrict Prosecutor's ability to send persuasive signals. Prosecutor has no incentive to increase Judge's payoff.

What is important to note is that our result of reduction in average informativeness does *not* have a negative effect on Judge accuracy.

4.2. Why two senders say less

Our final piece of analysis is a brief look at the mechanisms that drive the results presented above. *Why* do two senders provide less information than one sender if signals are costly, when two senders without costs provide substantially more information than a single sender in the absence of costs?

The central change in the strategic nature of the game that happens when we consider costly signals is that the game is no longer a zero-sum game. In the model without costs, the game is zero-sum because both senders' preferences are diametrically opposed over the binary action space of the Judge, and there is nothing else that is payoff relevant. Additionally, we construct the information environment in such a way that each sender can only increase the joint informativeness. Suppose a sender tries to strategically withhold information to persuade the judge in the preferred direction, thereby increasing their own expected payoff. In that case, this necessarily decreases the payoff of the other sender. However, each sender can never be worse off than their payoff under full revelation. As a result, full revelation by both senders is an equilibrium of this game, because everything is a best response to full revelation by the other sender, including also sending the fully revealing signal.

This zero-sum argument no longer holds if signals are costly. The unique best response to full revelation for each sender is the uninformative signal, because under full revelation of the other sender, they cannot change the outcome of the game and therefore maximize their payoff by choosing to incur zero costs. Moreover, full revelation is never best response to the uninformative signal, as an uninformative signal by the opponent is equivalent to a single sender setting, which we solved in section 3.1. This implies that under non-zero costs, there cannot be a pure strategy equilibrium that contains full revelation by either sender.

5. Discussion

5.1. Policy implications

Strategic information provision, as we describe it in our model, arises in various contexts. Apart from our example of a criminal trial, persuasion models can be well-suited to address settings such as advertising, where firms decide how much information about product quality to reveal; lobbying or political campaigning, where information about the effects of a specific policy is revealed; or the news industry, where media outlets provide selected information to frame a particular political issue in a certain light. It is worth reiterating that these examples typically involve *multiple* senders, underlining that our model can contribute to our understanding of such situations.

The literature tends to support the position that competition among multiple senders is beneficial for the receiver, as it leads to more information being revealed. In concrete terms, this could mean that when politicians consult lobbyists, they should always be required to hear from multiple interested groups to make more informed decisions.

Our results suggest that this conclusion should not be drawn without caution, or at least that the literature position might be overly optimistic. As we discussed above, the positive effect of increasing the number of senders on informativeness can rely on zero-sum competition among senders. In the introduction, we made the case that considering positive signal costs may be more reasonable in many settings. Based on our findings, it is therefore not clear to what extent we should expect gains in informativeness from increasing the number of senders in real-world situations.

Using these practical references, we aim to highlight that costly signals are a crucial additional assumption in various contexts, which lends relevance to the results obtained by including the cost function in our model. Ideally, a regulator would implement conditions in which information acquisition for the senders is costless, which would enable competition that possibly leads to full revelation.

Despite the results of our model, we are not comfortable drawing the policy recommendation to consult only singular information sources, especially in contexts such as political lobbying, from our findings.

Instead, we offer a more nuanced policy recommendation, first to consider whether the receiver can distinguish whether senders are likely to face costly signals in a specific setting. If a setting with costly information is identified, a way to restore informativeness could be to make targeted competition among the senders more difficult. For example, the receiver could keep the number and identities of competing senders secret from each other, such that each sender faces the illusion of a single-sender game, which we associate with more information in this context.

A caveat to our arguments is that pure informativeness may not always be the central criterion for welfare; therefore, it may not be the relevant variable on which to base our policy recommendations. To illustrate this point, recall that our theoretical evidence suggests that despite the loss in informativeness, the receiver's welfare did not decrease between the single-prosecutor and two-sender model under costly signals. Moreover, under costly signals, the receiver appears to be indifferent between receiving information from one sender, two senders, or, in fact, no sender at all (see Table 4.1).

5.2. Limitations

The elephant in the room is certainly that we do not provide a Nash equilibrium for the two-sender model with costly signals. As mentioned above, this caveat implies that we must exercise caution when comparing the outcomes of our ε-equilibrium strategy profiles to the equilibrium outcomes of other model specifications.

The way that we have set up the model leads to a fairly untractable set of payoff functions for the senders, in that their payoff depends on the discrete distribution of the Bayesian posterior, which results from the two binary signals. All of which is highly non-linear and discontinuous. What could be needed to provide a more adequate solution to the model is either a more sophisticated treatment of the algorithmic solution approach (for example, Hossain et al. (2024) provide an algorithmic solution to generalized multi-sender persuasion models using advanced neural networks), or possibly to impose additional restrictions that simplify the game without loss of generality.

Nonetheless, we defend that there are important lessons to be learned from our approximate solution. In the best-case scenario, a true equilibrium is indeed in the neighborhood, in which we characterize our ε–equilibrium and our findings are reasonably precise already. However, even in the worst-case, when the true equilibrium differs significantly from our guess, we do know for sure that full revelation is explicitly not an equilibrium of this game, as we argued in section 4.2. To that extent, we can at least refute the strong result that a receiver can harvest full revelation when facing two competing senders, even without pinning down a precise equilibrium. This insight alone is already valuable, as it points to the need for caution when considering the informational benefits of consulting multiple senders in situations of persuasion.

6. Conclusion

We began the thesis with the conventional wisdom that it is better to consult multiple sources of information. There certainly are many good examples of situations where this is the correct heuristic. Still, our model suggests otherwise, in that we should expect more senders to decrease informativeness under certain conditions.

To some extent, the findings obtained in this thesis are a success: we set out to challenge the result that more senders always reveal more information, and we found a plausible and relevant setting in which this might be the case.

On the other hand, our results are concerning in a world in which digital technologies have made the transmission of information increasingly easy. By implication, this also means that it is easy for additional senders to enter into a given situation of persuasion, which we would then associate with a reduction in the amount of provided information.

With signal costs, we describe one potential mechanism that can explain why two senders say less than one would. Of course, there are other relevant factors as well. For now, the conclusion we draw is that the effect of sender count on informativeness remains ambiguous.

We emphasize the importance of a thorough understanding of multi-sender persuasion, particularly regarding the impact of increased competition on the information disclosed. We hope to have contributed to this end and that the foundations laid out in this thesis will enable further insights in future work.

Bibliography

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

Bergemann, D. and Morris, S. (2016). Information design, bayesian persuasion, and bayes correlated equilibrium. *American Economic Review*, 106(5):586–91.

Boleslavsky, R. and Cotton, C. (2018). Limited capacity in project selection: competition through evidence production. *Economic Theory*, 65(2):385–421.

Degan, A. and Li, M. (2021). Persuasion with costly precision. *Economic Theory*, 72(3):869–908.

Gradwohl, R., Hahn, N., Hoefer, M., and Smorodinsky, R. (2022). Reaping the informational surplus in bayesian persuasion. *American Economic Journal: Microeconomics*, 14(4):296–317.

Hossain, S., Wang, T., Lin, T., Chen, Y., Parkes, D. C., and Xu, H. (2024). Multi-sender persuasion: A computational perspective.

Kamenica, E. (2019). Bayesian persuasion and information design. *Annual Review of Economics*, 11(Volume 11, 2019):249–272.

Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6):2590–2615.

Kamenica, E. and Gentzkow, M. (2014). Costly persuasion. *American Economic Review*, 104(5):457–462.

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

Kamenica, E. and Gentzkow, M. (2017b). Competition in persuasion. *Review of economic studies*, 84(1):1.

Koessler, F., Laclau, M., Renault, J., and Tomala, T. (2022). Long information design. *Theoretical Economics*, 17(2):883–927. ©2022 The Authors. Licensed under CC BY-NC 4.0.

Lemke, C. E. and Howson, Jr., J. T. (1964). Equilibrium points of bimatrix games. *Journal of the Society for Industrial and Applied Mathematics*, 12(2):413–423.

Li, F. and Norman, P. (2018). On bayesian persuasion with multiple senders. *Economics Letters*, 170:66–70.

Li, F. and Norman, P. (2021). Sequential persuasion. *Theoretical Economics*, 16(2):639–675. JEL classification: D82, D83.

Nguyen, A. and Tan, T. Y. (2021). Bayesian persuasion with costly messages. *Journal of Economic Theory*, 193:105212.

Ravindran, D. and Cui, Z. (2022). Competing persuaders in zero-sum games. `http://dx.doi.org/10.2139/ssrn.4241719`. Available at SSRN: `https://ssrn.com/abstract=4241719`.

Savani, R. and Turocy, T. L. (2025). Gambit: The package for computation in game theory, version 16.3.0. `https://www.gambit-project.org`.

Wu, W. (2023). Sequential bayesian persuasion. *Journal of Economic Theory*, 214:105763.

Wu, W. and Ye, B. (2023). Competition in persuasion: An experiment. *Games and Economic Behavior*, 138:72–89.

A. Appendix

A.1. Derivations

A.1.1. Covariance

Let

$$\mathrm{Cov}(s_P, s_D) = \mathbb{E}[s_P s_D] - \mathbb{E}[s_P]\mathbb{E}[s_D]$$

where

$$\mathbb{E}[s_i] = \sum_{\omega \in \Omega} \Pr(s_i|\omega)\Pr(\omega)$$

and

$$\mathbb{E}[s_P s_D] = \sum_{\omega \in \Omega} \Pr(s_i|\omega)\Pr(s_j|\omega)\Pr(\omega)$$

which yields:

$$\mathrm{Cov} = \mu_0(1-\mu_0)(\Pr(g_P|G)(\Pr(g_p|I)(\Pr(g_D|G)(\Pr(g_D|I)$$

Substituting signal probabilities for interim beliefs, we obtain:

$$\mathrm{Cov} = \frac{\lambda_P \lambda_D (\mu_P^H - \mu_0)(\mu_d^H - \mu_0)}{\mu_0(1-\mu_0)}$$

A.1.2. Linearity of cost function

We start from the original cost function:

$$C(\mu_i^H, \mu_i^L) = \lambda_i \cdot k(\mu_i^H - \mu_0)^2 + (1-\lambda_i) \cdot k(\mu_i^L - \mu_0)^2$$

with

$$\lambda_i = \frac{\mu_0 - \mu_i^L}{\mu_i^H - \mu_i^L}$$

Substituting λ_i:

$$C = k\left[\frac{\mu_0 - \mu_i^L}{\mu_i^H - \mu_i^L}(\mu_i^H - \mu_0)^2 + \left(1 - \frac{\mu_0 - \mu_i^L}{\mu_i^H - \mu_i^L}\right)(\mu_i^L - \mu_0)^2\right]$$

Letting $A := \mu_i^H - \mu_0$, $B := \mu_i^L - \mu_0$, we rewrite:

$$C = k\left[\frac{\mu_0 - \mu_i^L}{\mu_i^H - \mu_i^L}A^2 + \left(\frac{\mu_i^H - \mu_0}{\mu_i^H - \mu_i^L}\right)B^2\right]$$

Multiplying out:

$$C = k \cdot \frac{1}{\mu_i^H - \mu_i^L}\left[(\mu_0 - \mu_i^L)(\mu_i^H - \mu_0)^2 + (\mu_i^H - \mu_0)(\mu_i^L - \mu_0)^2\right]$$

Expanding both terms:

$$= k \cdot \frac{1}{\mu_i^H - \mu_i^L}\left[\mu_0(\mu_i^H - \mu_0)^2 - \mu_i^L(\mu_i^H - \mu_0)^2 + \mu_i^H(\mu_i^L - \mu_0)^2 - \mu_0(\mu_i^L - \mu_0)^2\right]$$

Now collecting like terms and simplifying leads to:

$$\begin{aligned} C(\mu_i^H, \mu_i^L) &= k\left(-\mu_0^2 + \mu_0\mu_i^H + \mu_0\mu_i^L - \mu_i^H\mu_i^L\right) \\ &= k(\mu_0 - \mu_i^L)(\mu_i^H - \mu_0) \end{aligned}$$

A.1.3. Entropy reduction and payoffs in other model specifications

We briefly denote the expected payoffs and expected entropy reduction for the two model specifications without costs and the single-prosecutor model with costly signals, to be illustrated in Table 4.1. The entropy of the prior is given as $H(\mu_0) = H(0.3) = 0.6109$.

Single Prosecutor without costs The single prosecutor setting is solved in the original persuasion model by Kamenica and Gentzkow (2011) with $\mu_P^{L^*} = 0, \mu_P^{H^*} = 0.5$, which implies $\lambda_P = 0.6$. We find the expected payoff as:

$$U_P = \lambda_P = 0.6$$

and calculate expected entropy reduction as

$$\mathbb{E}\Delta H = H(\mu_0) - [0.4H(0) + 0.6H(0.5)] \approx 0.1950$$

Two senders without costs As argued in Kamenica and Gentzkow (2017a) and Ravindran and Cui (2022), the two-sender model uniquely results in full revelation, which implies perfect receiver accuracy. Payoffs are therefore given as

$$U_P = \mu_0 = 0.3, U_D = 0.7$$

and since all prior entropy is reduced to zero,

$$\mathbb{E}\Delta H = H(\mu_0) - 0 \approx 0.6109$$

Single Prosecutor with costs We use the optimal strategy derived in 3.1, $\mu_P^{L^*} = 0.5 - \frac{1}{\sqrt{10}} \approx 0.18, \mu_P^{H^*} = 0.5$, which results in $\lambda_P \approx 0.37$ and yields

$$U_P = \lambda_P - c_P(0.18, 0.5) \approx 0.1351$$

and

$$\mathbb{E}\Delta H = H(\mu_0) - [0.63H(0.18) + 0.37H(0.5)] \approx 0.0540$$

A.2. Code

A.2.1. Model primitives and first equilibrium

```python

# This module constructs a discretized strategy grid, performs strict
# dominance elimination, and computes a mixed-strategy Nash equilibrium
# using the Gambit library.

import numpy as np
import pygambit as gb

# Model parameters
mu0  = 0.3      # Prior
k    = 10       # Cost parameter
step = 0.02     # Grid step size
tol  = 1e-8     # Tolerance for numerical comparisons

# ----------------- Core Model Functions -----------------
def continuous_cost(mu0, muL, muH, k, tol):
    """
    Define the cost for interim beliefs different from the prior.
```

```
    """
    if abs(muL - mu0) < tol:
        return 0.0
    return k * (-mu0**2 + mu0*muH + mu0*muL - muH*muL)


def posterior(a, b, mu0):
    """
    Define the posterior belief after a pair of interim beliefs a, b.
    Returns NaN if the denominator is (near) zero.
    """
    den = a*b - a*mu0 - b*mu0 + mu0
    if abs(den) < 1e-12:
        return float('nan')
    return (1.0 - mu0) * a * b / den


def compute_payoff(mu0, pL, pH, dL, dH, k, tol):
    """
    Given prosecutor strategy (pL, pH) and defender strategy (dL, dH),
    compute expected payoffs U_P and U_D.
    """
    # Probabilities of high signal for both players
    lp = 0.0 if abs(pH - pL) < tol else (mu0 - pL) / (pH - pL)
    ld = 0.0 if abs(dH - dL) < tol else (mu0 - dL) / (dH - dL)

    # Continuous costs of strategy
    cp = continuous_cost(mu0, pL, pH, k, tol)
    cd = continuous_cost(mu0, dL, dH, k, tol)

    # Covariance term
    cov = lp * ld * (pH - mu0) * (dH - mu0) / (mu0 * (1.0 - mu0))

    # Frequencies of possible posterior realizations
    fHH = lp*ld + cov
    fHL = lp*(1.0 - ld) - cov
    fLH = (1.0 - lp)*ld - cov
    fLL = (1.0 - lp)*(1.0 - ld) + cov

    # Save posteriors and frequencies
    posteriors = [
        posterior(pH, dH, mu0),
        posterior(pH, dL, mu0),
        posterior(pL, dH, mu0),
```

```python
            posterior(pL, dL, mu0),
        ]
        freqs = [fHH, fHL, fLH, fLL]

        U_P = U_D = 0.0
        # Calculate payoffs as sum of frequencies of posterior realizations
        for post, freq in zip(posteriors, freqs):
            if post >= 0.5:
                U_P += freq
            elif not np.isnan(post):
                U_D += freq

        return U_P - cp, U_D - cd

# --------------- Strategy Construction & Reduction ---------------
def build_strategy_grid(mu0, step):
    """
    Construct  discrete list of pure strategies.
    """
    grid = []
    for muL in np.arange(0, mu0, step):
        for muH in np.arange(mu0 + step, 1 + tol, step):
            if muL < mu0 < muH:
                grid.append((round(muL, 3), round(muH, 3)))

    grid.append((mu0, round(mu0 + step, 3)))
    return sorted(set(grid))


def eliminate_dominated(pros_set, def_set, mu0, k, tol):
    """
    Remove strictly dominated strategies from both sets by comparing payoffs.
    Returns reduced sets and lists of eliminated strategies.
    """
    pros_arr = np.array(pros_set)
    def_arr  = np.array(def_set)

    # Compute payoff matrices for all pairs
    Up = np.empty((len(pros_arr), len(def_arr)))
    Ud = np.empty_like(Up)
    for i, (pL,pH) in enumerate(pros_arr):
        for j, (dL,dH) in enumerate(def_arr):
            Up[i,j], Ud[i,j] = compute_payoff(mu0, pL, pH, dL, dH, k, tol)

```

```python
    # Identify dominated strategies
    dom_p = {pros_set[i] for i in range(Up.shape[0]) if any((Up[t,:] > Up[i,:]).all() for t in range(Up.shape[0]) if t!=i)}
    dom_d = {def_set[j] for j in range(Ud.shape[1]) if any((Ud[:,e] > Ud[:,j]).all() for e in range(Ud.shape[1]) if e!=j)}

    new_p = [s for s in pros_set if s not in dom_p]
    new_d = [s for s in def_set  if s not in dom_d]
    return new_p, new_d, dom_p, dom_d

def iterate_elimination(pros_set, def_set, mu0, k, tol):
    """
    Iteratively apply eliminate_dominated until no strategies are removed.
    Returns surviving strategy sets.
    """
    pros, defs = pros_set.copy(), def_set.copy()
    iteration = 0
    while True:
        iteration += 1
        new_p, new_d, dom_p, dom_d = eliminate_dominated(pros, defs, mu0, k, tol)
        if not dom_p and not dom_d:
            break
        pros, defs = new_p, new_d
    return pros, defs

# --------------- Equilibrium Computation with Gambit ---------------
def solve_equilibrium(A, B, seed=0):
    """
    Use Gambit's Lemke-Howson solver to find a Nash equilibrium in zero-sum game defined
    by payoff matrices A (prosecutor) and B (defender).
    """
    # Ensure non-negative payoffs
    shift = -min(A.min(), B.min(), 0.0)
    A0, B0 = A + shift, B + shift

    game = gb.Game.from_arrays(A0, B0)
    solution = gb.nash.lcp_solve(game, rational=False)
    eqs = solution.equilibria if hasattr(solution, 'equilibria') else list(solution)
    if not eqs:
        raise RuntimeError("No equilibrium found after shifting payoffs.")
    return eqs[0], shift
```

A.2.2. Simulation of continuous stratgy profile

```python
#  Monte Carlo simulation setup
h = 0.02        # width of grid cell for noise
delta = 0.001   # half-thickness for thin-strip sampling

import math, time

# Extract equilibrium support strategies and probabilities from discrete equilibrium
P_player, D_player = game.players
pros_probs = np.array([eq_profile[P_player][s] for s in pros_surv], float)
defs_probs = np.array([eq_profile[D_player][s] for s in def_surv], float)
pros_support = np.column_stack((pros_surv, pros_probs))
defs_support = np.column_stack((def_surv, defs_probs))

# Helper: binary entropy function
def entropy(p):
    """Binary entropy; returns 0 at p=0 or p=1."""
    if p <= 0.0 or p >= 1.0:
        return 0.0
    return - (p * math.log(p) + (1 - p) * math.log(1 - p))
prior_ent = entropy(mu0)

# Helper: convert interim beliefs to signal probabilities
def compute_signal_probs(muL, muH):
    if abs(muH - muL) < tol:
        return mu0, mu0
    alpha = (muH * (mu0 - muL)) / (mu0 * (muH - muL))
    beta = ((mu0 - muL) * (1 - muH)) / ((1 - mu0) * (muH - muL))
    return np.clip(alpha, 0.0, 1.0), np.clip(beta, 0.0, 1.0)

# Helper: Bayes' update given signals
def bayes_posterior(alpha_P, beta_P, alpha_D, beta_D, sigP, sigD):
    L_G = (alpha_P if sigP == 'g' else 1 - alpha_P) * \
          (alpha_D if sigD == 'g' else 1 - alpha_D)
    L_I = (beta_P if sigP == 'g' else 1 - beta_P) * \
          (beta_D if sigD == 'g' else 1 - beta_D)
    denom = mu0 * L_G + (1 - mu0) * L_I
    return mu0 if denom < 1e-15 else (mu0 * L_G) / denom

# Helper: sample a pure strategy from the continuous mixture
```

```
def sample_sender(support, n, role, rng):
    centers_L, centers_H, probs = support[:, 0], support[:, 1], support[:, 2]
    idx = rng.choice(len(probs), size=n, p=probs)
    muL_c, muH_c = centers_L[idx], centers_H[idx]
    if role.upper() == 'P':
        # vertical strip:
        muL = muL_c + rng.uniform(-delta, delta, size=n)
        muH = muH_c + rng.uniform(-h/2, h/2, size=n)
    else:
        # horizontal strip:
        muL = muL_c + rng.uniform(-h/2, h/2, size=n)
        muH = muH_c + rng.uniform(-delta, delta, size=n)
    return np.clip(muL, 0.0, mu0 - 1e-12), np.clip(muH, mu0 + 1e-12, 1.0)

# Helper: label joint signal profiles
def encode_profile(sigP, sigD):
    return (sigP == 'g') * 2 + (sigD == 'g')

# Initialize simulation
N_SIM = 500_000
rng = np.random.default_rng(12325)
state_arr = rng.random(N_SIM) < mu0          # true states
actions = np.empty(N_SIM, int)              # decision outcomes
uP_arr = np.empty(N_SIM)                    # payoffs for P
uD_arr = np.empty(N_SIM)                    # payoffs for D
post_arr = np.empty(N_SIM)                  # posterior beliefs
ent_red = np.empty(N_SIM)                   # entropy reductions

# Main Monte Carlo loop
t0 = time.time()
for t in range(N_SIM):
    s = 1 if state_arr[t] else 0                # binary state
    pL, pH = sample_sender(pros_support, 1, 'P', rng) # draw prosecutor strategy
    dL, dH = sample_sender(defs_support, 1, 'D', rng) # draw defender strategy
    alpha_P, beta_P = compute_signal_probs(pL, pH) # convert strategies to signals
    alpha_D, beta_D = compute_signal_probs(dL, dH)
    # generate signal realizations based on state and signal probs
    sigP = 'g' if rng.random() < (alpha_P if s else beta_P) else 'i'
    sigD = 'g' if rng.random() < (alpha_D if s else beta_D) else 'i'
    p_post = bayes_posterior(alpha_P, beta_P, alpha_D, beta_D, sigP, sigD) # posterior
    ent_red[t] = prior_ent - entropy(p_post) # entropy reduction
    a = 1 if p_post >= 0.5 else 0               # action: convict if posterior >= 0.5
    uP_arr[t], uD_arr[t] = compute_payoff(mu0, pL, pH, dL, dH, k, tol)
    actions[t], post_arr[t] = a, p_post
```

```
t1 = time.time()

# Summary of results
receiver_acc = (actions == state_arr).mean()   # accuracy of decision-maker
print(f"Simulation done in {t1 - t0:.2f}s (N={N_SIM:,})")
print(f"Mean U_P: {uP_arr.mean():.4f}, Mean U_D: {uD_arr.mean():.4f}")
print(f"Receiver accuracy: {receiver_acc:.4f}")
print(f"Entropy reduction: {ent_red.mean():.4f}")

```

A.2.3. Simulation for Deviations

```
#!/usr/bin/env python
# coding: utf-8

# In[ ]:


# Use Monte Carlo results (uP_arr, uD_arr) as baseline equilibrium payoffs
Up_eq = uP_arr.mean()   # Equilibrium payoff for Prosecutor
Ud_eq = uD_arr.mean()   # Equilibrium payoff for Defender

# Vectorized payoff evaluation for deviation candidates
payoff_vec = np.vectorize(
    lambda pL, pH, dL, dH: compute_payoff(mu0, pL, pH, dL, dH, k, tol),
    otypes=[float, float]
)

# Find best pure-strategy deviation given opponent's support
import time
N_dev = 10_000    # number of candidate deviations to test per player
M     = 10_000    # number of mixed strategy samples per deviation candidate
rng   = np.random.default_rng(12325)

def best_deviation(role):
    """
    Search for the pure strategy (muL, muH) that maximizes excess payoff
    over equilibrium for the specified role ('P' or 'D').
    """
    if role == 'P':
        # Sample defender strategies from equilibrium support
        dL_s, dH_s = sample_sender(defs_support, M, 'D', rng)
        # Generate candidate deviations across fine fine-tuned deviation area
```

```
        dev_L = rng.uniform(0.2, mu0,   N_dev)
        dev_H = rng.uniform(mu0, 0.5, N_dev)
        Up_dev, _ = payoff_vec(dev_L[:, None], dev_H[:, None], dL_s[None, :], dH_s[None, :])
        excess = Up_dev.mean(axis=1) - Up_eq
    else:
        # Sample prosecutor strategies
        pL_s, pH_s = sample_sender(pros_support, M, 'P', rng)
        # Generate candidate deviations across fine fine-tuned deviation area
        dev_L = rng.uniform(0.2, mu0,   N_dev)
        dev_H = rng.uniform(mu0, 0.5, N_dev)
        _, Ud_dev = payoff_vec(pL_s[None, :], pH_s[None, :], dev_L[:, None], dev_H[:, None])
        excess = Ud_dev.mean(axis=1) - Ud_eq
    # Identify best deviation
    idx = np.argmax(excess)
    return dev_L[idx], dev_H[idx], excess[idx]

# Execute deviation search
tic = time.time()
best_P_L, best_P_H, eps_P = best_deviation('P')
best_D_L, best_D_H, eps_D = best_deviation('D')
toc = time.time()

# Report search results
print("
Deviation Search Results:")
print(f" Runtime: {toc - tic:.1f}s (N_dev={N_dev:,}, M={M:,})")
print(f" Prosecutor best deviation: mu_L={best_P_L:.3f}, mu_H={best_P_H:.3f}, _P={eps_P:.6f}")
print(f" Defender best deviation:   mu_L={best_D_L:.3f}, mu_H={best_D_H:.3f}, _D={eps_D:.6f}")
```