Money Burning Improves Mediated Communication*

Yi Liu[†] Yang Yu[‡]

November 16, 2024

Abstract

This paper explores the problem of mediated communication enhanced by money-burning tactics for commitment power. In our model, the sender has state-independent preferences and can design a communication mechanism that both transmits messages and burns money. We characterize the sender's maximum equilibrium payoff, which has clear geometric interpretations and is linked to two types of robust Bayesian persuasion. We demonstrate that, generically, the money-burning tactic *strictly* improves the sender's payoff for almost all prior beliefs where commitment is valuable for the sender. Furthermore, our communication model directly applies to Web 3.0 communities, clarifying the commitment value within these contexts.

Keywords: Mediated Communication, Money Burning, Mechanism Design, Cheap Talk, Bayesian Persuasion, Commitment;

JEL Classification: D82, D83.

*We thank Nemanja Antic, Roberto Corrao, Yifan Dai, Piotr Dworczak, Françoise Forges, Alexander Jakobsen, Zhonghong Kuang, Xiao Lin, Harry Pei, , Fan Wu, Xingye Wu, Junjie Zhou and the audiences at Tsinghua University-Peking University NCER-CCER the third China Economic Seminar for their helpful comments and discussions. This paper was circulated with the title "Mediated Information Design with Money Burning for Commitment Power."

[†]Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing, China, 100084 (liu-yi22@mails.tsinghua.edu.cn).

[‡]School of Economics and Management, China University of Petroleum-Beijing, Beijing, China (yangyu@cup.edu.cn).

1 Introduction

What is the best Sender (he) can do *by himself* in strategically communicating with uninformed Receiver (she)? By considering different levels of the Sender's commitment, various communication protocols emerge. When Sender has no commitment, Crawford and Sobel [1982] introduce the cheap talk (CT) model, in which Sender can freely send unverifiable messages to Receiver. In contrast, when Sender has full commitment, Kamenica and Gentzkow [2011] present the Bayesian persuasion (BP) model, allowing the Sender to ex ante commit to a communication strategy or an experiment that is verifiable by both parties. However, obtaining such commitment in real-world contexts is challenging, and the unverifiable nature of cheap talk renders it inefficient in many scenarios. As an intermediate level of commitment, where Sender can only commit to the message-generating process based on his report but not to the entire communication strategy[1], Salamanca [2021] examines the optimal Sender's equilibrium payoff within the mediated communication (MD) model introduced by Myerson [1982], Forges [1986]. Nonetheless, is the mediated communication protocol optimal for Sender, given that he can commit to the message-generating process?

Here, we propose a novel communication protocol called *mediated communication with money-burning mechanism* (or MDMB for short), whereby Sender employs the money-burning tactic to obtain credibility. The MDMB includes two parts pre-determined by Sender: the information-transmission rule for the mediator sending the message and the extent of money-burning based on the Sender's report. In this protocol, Sender can independently implement the money-burning tactic through the mediator, without requiring additional assistance from other parties, given the current level of the Sender's commitment power.

We investigate our proposed communication protocol under a substantive assumption: Sender has state-independent preferences over the Receiver's actions. This transparent-motives assumption simplifies the analysis while retaining substantial real-world economic applications Chakraborty and Harbaugh [2010], Lipnowski and Ravid [2020], Lipnowski et al. [2022].

We also found that MDMB has been applied as the communication paradigm in

[1]Alternatively, we can consider a scenario where there is a trusted mediator. In this case, Sender communicates directly with the mediator, who then relays the information to the Receiver according to a predetermined rule established by Sender.

the emerging Web 3.0 economy as proposed by Drakopoulos et al. [2023]. For instance Shaker et al. [2021], some Web 3.0 financial companies sell their products to consumers through smart contracts. Those companies input the risk information into the smart contracts. The smart contracts generate the risk assessment results with randomness to the consumers according to the pre-decided and transparent algorithms. Here, those companies are Sender while consumers are Receivers. The smart contracts are the mediator. The gas fee or token transferred from the financial companies to the consumers is the money-burning tactics. In all Web 3.0 business practices, the transparent and auto-processed algorithms simultaneously play the role of a trustworthy mediator for enforcing the pre-committed message-design protocol and enabling the money burning by gas fee. Therefore, MDMB is a general communication protocol in the Web 3.0 economy.

In this paper, we primarily address two questions. First, what is the Sender's optimal equilibrium payoff when utilizing MDMBs, and what is the corresponding optimal design? Second, and more importantly, under what conditions can money burning *strictly* improve the boundaries of mediated communication?[2] This question is crucial for justifying the use of the money-burning tactic.

Analyzing the MDMB is technically challenging. The combination of mechanism design and information transmission significantly increases the difficulty of identifying and analyzing the optimal MDMB. The burned money may contain information about the Sender's type, allowing the Receiver to select her action accordingly. Consequently, the outcomes of MDMBs can influence the Receiver's actions in ways that the Sender cannot commit to, which violates the full-commitment condition necessary for applying the revelation principle Bester and Strausz [2001]. Thus, the revelation principle cannot be used to simplify the non-convex optimal MDMB mechanism problem restricted by the equilibrium constraints.

We develop a new revelation principle for MDMB that decomposes it into a separable sequential process. According to this principle, Sender first designs the message and then determines the amount of money to be burned based on that message. This design ensures that the burned money contains no additional information about the Sender's type beyond what is conveyed in the message. It also allows the Receiver to directly check the burned amount based on the received messages, which in turn reduces the commitment

[2] It is obvious that money burning weakly improves mediated communication since Sender can always choose not to burn money.

requirement for burning money. This decomposition enables us to employ a belief-based approach, transforming the non-convex optimal MDMB problem constrained by equilibrium conditions into a simpler optimization problem subject to incentive-compatible and Bayes-plausible constraints. The revelation principle for MDMB introduces a novel framework for mechanism design under limited commitment, complementing the revelation principle for the dynamic mechanism selection game developed in Doval and Skreta [2022].

Using the decomposition structure in the revelation principle for MDMB, we characterize the Sender's maximum equilibrium payoff, known as the value of MDMB, along with the associated optimal MDMB design. Proposition 3 geometrically illustrates that the value of MDMB is the minimum value among all concavification values of the Sender's subjective payoff functions[3]. This geometric analysis reveals how money burning enhances commitment power. In the optimal design, the amount of money burned varies with the messages that truthfully reveal the Sender's type with a small unconditional probability, ensuring that Sender can only achieve the minimum interim payoff for all of his types. Consequently, Sender has no incentive to deviate from truthfully reporting his type.

Theorem 1 demonstrates that, generically, Sender has a strictly higher payoff after employing the money-burning tactic in MD for almost all prior beliefs when the commitment is valuable. Although we now have a clear characterization of the value of MDMB and the value of MD[4], we cannot directly compare these two values by examining the sets of the worst convex combinations in Proposition 3 and the worst affine combinations in the characterization of the value of MD. This limitation arises because there are limited insights into the sets of the worst convex and affine combinations in general. Therefore, we adopt an alternative approach to obtain Theorem 1.

We first identify a necessary condition under which money burning cannot improve MD, and then we employ a constructive approach to demonstrate that this necessary condition generically does not hold. To this end, we define a *generic* condition on the Receiver's payoff functions, where any action that is Receiver-optimal under certain beliefs is uniquely optimal under other beliefs with the same support. Proposition 7 establishes

[3] The convex combination (expectation according to the subjective belief) of truth-adjusted welfare functions defined in Doval and Smolin [2024].

[4] See Corrao and Dai [2023], Salamanca [2021]; the value of MD here is the minimum value among all concavification values of the Sender's virtual value functions which is an affine combination of adjusted welfare functions.

that for all Receiver's payoff functions satisfying this generic condition, the value of MDMB equals the value of MD if and only if the value of CT is also equal to this value. Therefore, we only need to show that the value of MDMB is greater than the value of CT, which is much easier than directly comparing MDMB and MD. We then utilize the characterization of the value of MDMB in Proposition 3, demonstrating that for almost all priors, the value of MDMB is strictly larger than that of CT (see Proposition 8).

We also notice that the above conclusion implies the tie between the MDMB and the robust Bayesian persuasion problems. We discover that the value of MDMB is also equal to the maximum payoff of the cautious Sender with full commitment power[5]. We also demonstrate that the value of MDMB is equal to the payoff of Sender facing the worst subjective prior in BP with heterogeneous beliefs Alonso and Câmara [2016]. The above discoveries suggest that the analysis of robust BP can be used to facilitate the MDMB analysis and also justifies the assumptions on min-max utility in the robust BP problem.

Since MDMB serves as the general model capturing the communication process in the Web 3.0 economy when commitment power is absent, the gap between MDMB and BP highlights the necessity of commitment and is referred to as the refined value of commitment in Web 3.0 economy. Proposition 9 demonstrates that commitment in Web 3.0 is valuable if and only if commitment in the conventional economy is valuable. However, Theorem 1 indicates that the refined value of commitment in Web 3.0 is less than the value of commitment in the conventional economy in almost all scenarios. Therefore, the algorithms used in Web 3.0 partially mitigate the losses associated with the absence of commitment power.

As extensions, we further explored the MDMB that is limited by a finite budget. Proposition 10 shows that the value of MDMB limited by budget constraint is the worst concavification value of Sender's generalized subjective payoff function which is the affine combination of Sender's generalized adjusted payoff functions. The generalized adjusted payoff functions further manifest Sender's optimal MDMB design. The message of an optimal MDMB can be classified into two groups. One group of messages is not associated with money burning and only functions for persuasion. The other group of messages is associated with money burning and used for obtaining credibility. In this group, the amount of money burned is maximized with the consideration of all the constraints.

[5] Cautious Sender is one who maximizes his minimum payoff across types Doval and Smolin [2021]

1.1 Related Literature

This paper proposes a novel communication protocol where adopting the money-burning tactic enhances the credibility in a limited commitment environment. Our work contributes to the literature that studies communication protocols under various degrees of commitment power. In addition to the literature mentioned earlier, Min [2021], Lipnowski et al. [2022] concentrate on the case where Sender's commitment power has a Bernoulli distribution on full commitment and no commitment; and Lin and Liu [2024] examine the situation where Sender cannot commit to the message-generating process but he can commit to the marginal distribution of types and messages. Bergemann and Morris [2019] summarize information design problems involving persuasion and mediation. Furthermore, our paper is closely related to studies that analyze the effect of commitment on information design under various communication protocols. Fréchette et al. [2022] investigate the effect of communication with different levels of commitment power through experimental methods. Additionally, Corrao and Dai [2023] comparatively analyze different communication protocols at various levels of commitment power, but they do not take money burning into account.

The main contribution of this paper is to extend the domain of mediated communication problems. Previous research, such as Salamanca [2021], illustrates the optimal equilibrium payoff of Sender through mediated communication without burning money. Drakopoulos et al. [2023] establish a blockchain system as a mediator, demonstrating that designing costly messages can improve MD under transparent motives, but they do not identify the optimal Sender's communication efficiency in general as we do and they also do not identify the condition where costly message improves MD. Additionally, several studies, including Goltsman et al. [2009], Ivanov [2014], have identified the optimal mediation plan for Receiver. Furthermore, Ivanov [2014] compare the outcomes of mediated communication and cheap talk.

Our paper is also related to the literature on communication with transfers. Some studies discuss cheap talk involving money burning, Austen-Smith and Banks [2000], Kartik [2007], Karamychev and Visser [2017], noting that Sender nearly cannot improve the credibility of cheap talk by money burning tactic, and chooses not to burn money even in state-independent preferences environments. However, in our work, money-burning mechanism plays a key role in enhancing Sender's commitment power and thus obtaining better communication efficiency. Kolotilin and Li [2021] investigate the application of monetary transfers in repeated cheap talk settings, while Sadakane [2023] examines a model featuring repeated cheap talk games with monetary transfers from Receiver to Sender. This latter study observes that the equilibrium set in such settings is larger than that of the original long-term cheap talk setting. Corrao [2023] analyzes the mediation market and characterizes the information and market outcomes of the revenue-maximizing mediator and the Sender-optimal mediator. Additionally, several studies focus on Bayesian persuasion involving transferable utility and the cost of information, such as Li and Shi [2017], Bergemann et al. [2018]. Dughmi et al. [2019] explore the case where Sender can enter into contracts prior to persuasion, while Perez-Richet and Skreta [2022] investigate the Receiver-optimal experiment under the condition that Sender can costly falsify his private type.

Another important category of literature related to us is about mechanism design with limited commitment. Liu and Wu [2024] examine the implementation problem in general outcome-contingent settings, which is a more generalized context than ours. Bester and Strausz [2001] show that the revelation principle fails to hold in a limited commitment environment, where the principal cannot fully commit to the outcome induced by the mechanism. Doval and Skreta [2022] provide the general revelation principle for limited commitment mechanism design, where the joint design of information and mechanism can be separated into two steps: first, design the information, and second, design the mechanism based on the information.

2 Model

In Section 2.1, we develop the basic model of the Sender-Receiver game, and in Section 2.2, we introduce the methodology for simplifying the Sender's programming problem.

2.1 Basic Setup

Primitives. We consider a basic game with two players: Sender (he) and Receiver (she). Sender has the private information θ, which denotes his type and belongs to a finite set Θ. The type θ is drawn according to a prior distribution $\mu_0 \in \Delta(\Theta)$, which is a common knowledge. Receiver can choose an action a from a finite set A, which determines the payoffs of both players depend. Receiver's payoff also depends on the Sender's type $\theta \in \Theta$. Sender's value function of Receiver's action is $v(\cdot) : A \to \mathbb{R}$ and the Receiver's value function is $u(\cdot, \cdot) : A \times \Theta \to \mathbb{R}$. Both players are risk-neutral and fully rational.

Communication with money-burning mechanisms. Before the game, Sender commits a mediated communication with money-burning mechanism (or MDMB). An MDMB consists of an input set M, an output message set S, and a corresponding mechanism $\phi : M \to \Delta(S \times \mathbb{R}_{\geq 0})$. The MDMB prescribes how Sender designs the message and determines the money-burning amount according to his private input. Here, we restrict attention to the case that M, S, and the support set of ϕ are all finite.

The Sender-Receiver game. The timeline of the MDMB as an extensive-form game is summarized as below.

Stage 1. Sender commits to the MDMB (M, S, ϕ) with a mediator.

Stage 2. The Sender's type θ is revealed to him according to the prior distribution μ_0. And then Sender sends an input message $m \in M$ to the mediator.

Stage 3. The mediator sends an output message $s \in S$ to Receiver and burns $t \geq 0$ money from the Sender's account, with probability $\phi(s, t|m)$.

Stage 4. Receiver observes the money burnt by Sender and the message m, updates her belief, and chooses an action $a \in A$.

Stage 5. Receiver gets payoff $u(a, \theta)$ and Sender gets payoff $v(a) - t$.

To analyze the optimal MDMB for Sender, we first apply backward induction to examine the Perfect Bayesian equilibrium of sub-game spanning stage 2 through stage 5, denoted as $\mathcal{G}_{(M,S,\phi)}(\mu_0)$.

Beliefs and strategies. The Sender's strategy in $\mathcal{G}_{(M,S,\phi)}(\mu_0)$ prescribes a transition probability $\sigma : \Theta \to \Delta(M)$. Here, $\sigma(\theta)$ denotes the probability distribution of input messages from Sender to Receiver when the Sender's type is θ. Receiver is only informed about the output message and the money burning amount. Hence, the output message s and the money burning amount t together formulate the information set of Receiver. For each information set (s, t), the Receiver's strategy prescribes a transition probability

$\alpha : S \times \mathbb{R}_{\geq 0} \to \Delta(A)$. Here, $\alpha(s,t)$ denotes the probability distribution of the actions responding to (s,t). In each information set (s,t), Receiver must form a belief $\mu : S \times \mathbb{R}_{\geq 0} \to \Delta(\Theta)$, where $\mu(s,t)$ denotes the probability distribution of Sender's types given information set (s,t). We call the triple (σ, α, μ) an assessment.

Equilibrium. In this paper, we use Perfect Bayesian equilibrium (henceforth, PBE) as the solution concept of game $\mathcal{G}_{(M,S,\phi)}(\mu_0)$. We denote the set of PBE of game $\mathcal{G}_{(M,S,\phi)}(\mu_0)$ as $\mathcal{E}[\mathcal{G}_{(M,S,\phi)}(\mu_0)]$. An assessment (σ, α, μ) is a PBE if it is sequentially rational and the belief μ satisfies Bayes' rule where possible. Formally, an assessment $(\sigma^*, \alpha^*, \mu^*)$ is a PBE, $(\sigma^*, \alpha^*, \mu^*) \in \mathcal{E}[\mathcal{G}_{(M,S,\phi)}(\mu_0)]$, if it satisfies following three conditions:

Sender's optimality: for any $\theta \in \Theta$,

$$\sigma^*(\theta) \in \arg \max_{\sigma(\theta) \in \Delta(M)} \sum_{m \in M, s \in S, t \geq 0, a \in A} \sigma(m|\theta)\phi(s,t|m)\alpha^*(a|s,t)(v(a) - t). \tag{1}$$

Receiver's optimality: for any $s \in S, t \geq 0$,

$$\alpha^*(s,t) \in \arg \max_{\alpha(s,t) \in \Delta(a)} \sum_{\theta \in \Theta, a \in A} \mu^*(\theta|s,t)\alpha(a|s,t)u(a,\theta). \tag{2}$$

Bayes updating: for any $s \in S, t \geq 0$ and $\theta \in \Theta$,

$$\mu^*(\theta|s,t) \sum_{\theta' \in \Theta, m \in M} \mu_0(\theta')\sigma^*(m|\theta')\phi(s,t|m) = \mu_0(\theta) \sum_{m \in M} \sigma^*(m|\theta)\phi(s,t|m). \tag{3}$$

Communication efficiency. $\mathcal{E}[\mathcal{G}_{(M,S,\phi)}(\mu_0)]$ enables us to analyze the optimal MDMB in stage 1. Given that Sender seeks to maximize his ex ante expected payoff across all possible PBEs, we can formulate the Sender's MDMB optimization problem as follows:

$$\begin{aligned} &\sup_{M,S,\phi} \sum_{\theta \in \Theta} \mu_0(\theta) \sum_{m \in M, s \in S, t \geq 0, a \in A} \sigma^*(m|\theta)\phi(s,t|m)\alpha^*(a|s,t)(v(a) - t) \\ &s.t. (\sigma^*, \alpha^*, \mu^*) \in \mathcal{E}[\mathcal{G}_{(M,S,\phi)}(\mu_0)]. \end{aligned} \tag{4}$$

The value of Equation 4 is referred to as the communication efficiency of MDMB, or alternatively, *the value of MDMB*, denoted as $\mathcal{V}^*(\mu_0)$.

2.2 Simplifying the Problem

The optimal MDMB design problem Equation 4 is complex due to its non-convex and equilibrium-selection complexity. However, the revelation principle developed by Myerson [1982], Forges [1986] cannot be applied to simplify the optimal MDMB problem. The amount of money burning can contain information about the Sender's type. Thus, the MDMB design will influence the Receiver's action that Sender cannot commit to and thus violate the condition of applying the revelation principle Bester and Strausz [2001].

In this section, we develop a new technique to simplify the MDMB. Inspired by Doval and Skreta [2022], we apply the method of canonical mechanisms and canonical assessments to develop a new revelation principle for MDMB. We then use the revelation principle for MDMB and belief-based approach to convert the original complex problem into an optimization problem under incentive-compatible and Bayes-plausible constraints.

Here, we first formally define the canonical MDMBs and canonical assessments under which we can calculate the value of MDMB without loss of generality.

Definition 1 (Canonical MDMBs)**.** *An MDMB is canonical if $M = \Theta$, $S = \Delta(\Theta)$, and there exists a signaling scheme $\pi : \Theta \to \Delta(\Delta(\Theta))$ and a deterministic function $x : \Delta(\Theta) \to \mathbb{R}_{\geq 0}$ such that π satisfies the Bayes updating condition*[6] *and $\phi(u, x(u)|\theta) = \pi(u|\theta)$ for all $\theta \in \Theta$ and $u \in \boldsymbol{supp}\{\pi(\theta)\}$.*

In canonical MDMBs, the input sets to a MDMB are the type sets while the output sets are the sets of distributions of types. The output message of the canonical MDMB contains *all* information transmitted to the Receiver, and the amount of money burning in the canonical MDMB does not provide any extra piece of information about the Sender's type beyond that contained in the output message. Hence, ϕ in a canonical MDMB can be decomposed into two parts. The first part is a *signaling scheme* π, and the second part is a *money-burning scheme* x which is contingent on the output message. This decomposition has a similar structure to the revelation principle in Doval and Skreta [2022]. Henceforth, we use (π, x) to refer to a canonical MDMB.

In a canonical MDMB, the canonical assessment ensures that the Sender's strategy is truthful-telling and the Receiver's posterior belief coincides with the output message.

[6] $u(\theta) \sum_{\theta' \in \Theta} \mu_0(\theta')\pi(u|\theta') = \mu_0(\theta)\pi(u|\theta)$ for all $u \in \Delta(\Theta)$.

Definition 2 (Canonical assessments)**.** *For a canonical MDMB, an assessment* (σ, α, μ) *is canonical if* $\sigma(\theta|\theta) = 1$ *and* $\mu(u, x(u)) = u$ *for any* $u \in \textbf{\textit{supp}}\{\pi(\theta)\}$*.*

The following proposition explains that every MDMB has a corresponding canonical MDMB that maintains the same expected payoff of Sender. This proposition allows us to focus only on the canonical MDMBs and the associated canonical assessment without loss of generality as depicted in Figure 1.

Proposition 1. *For any MDMB* (M, S, ϕ) *and* $(\sigma, \alpha, \mu) \in \mathcal{E}[\mathcal{G}_{M,S,\phi}(\mu_0)]$*, there exists a canonical MDMB* (π, x) *and a canonical assessment* $(\sigma^*, \alpha^*, \mu^*) \in \mathcal{E}[\mathcal{G}_{(\pi,x)}(\mu_0)]$ *such that the expected payoffs of Sender in both assessments are the same.*

Proof. This proof is relegated to Appendix B. □

$$\underset{\theta}{\text{Sender}}\ M \xrightarrow{\phi(\cdot|m)} S \times \mathbb{R}_{\geq 0}\ \text{Receiver} \qquad \Leftrightarrow \qquad \underset{\theta}{\text{Sender}}\ \Theta \xrightarrow{\pi(\cdot|\theta)} \Delta(\Theta) \xrightarrow{x(\cdot)} \mathbb{R}_{\geq 0}\ \text{Receiver}$$

Figure 1: Revelation principle

In the rest of this section, we explain how to apply Proposition 1 to simplify the optimal MDMB problem. We first define a sequence of necessary concepts related to the belief-based approach. Then, we explain how to apply Proposition 1 and the belief approach to 1) convert Sender's optimality constraints to incentive-compatible constraints and 2) convert the Bayes updating and Receiver's optimality constraints to the Bayes plausible constraints.

We first define the Receiver-optimal set as the set of Receiver's best responses according to posterior belief $\mu \in \Delta(\Theta)$ denoted it as

$$RO(\mu) \triangleq \big\{\alpha \in \Delta(A) \big| \textbf{supp}\{\alpha\} \subset \arg\max_{a' \in A} \sum_{\theta \in \Theta} \mu(\theta) u(a', \theta)\big\}.$$

According to Receiver's posterior belief, we further define Sender's *belief-value correspondence* as $\mathbb{V} : \Delta(\Theta) \rightrightarrows \mathbb{R}$, where $\mathbb{V}$ is the collection of all possible ex-post signaling payoffs for Sender and expressed as

$$\mathbb{V}(\mu) \triangleq \big\{q \big| \exists \alpha \in RO(\mu), q = \sum_{a \in A} \alpha(a) v(a)\big\}.$$

Furthermore, let $p \in \Delta(\Delta(\Theta))$ denote the distribution over the Receiver's posterior belief induced by a PBE canonical assessment. Then, we can define the Bayesian plausible set associated with the prior μ_0 according to Kamenica and Gentzkow [2011] as

$$BP(\mu_0) \triangleq \{p \in \Delta(\Delta(\Theta)) \big| \int_\mu \mu dp(\mu) = \mu_0\}.$$

Then, we in following proposition transform the equilibrium constraints in the equilibrium selection problem of Equation 4 into incentive compatibility, obedience, and Bayesian plausibility constraints.

Proposition 2. *A distribution over the Receiver's posterior belief p and the Sender's ex-post signaling payoff $V : \Delta(\Theta) \to \mathbb{R}$ in terms of the posterior belief can be induced by a PBE canonical assessment of a canonical MDMB if and only if p and V fulfill the following conditions:*

Incentive compatibility: for any $\theta, \theta' \in \Theta$,

$$\int_\mu \left(\frac{\mu(\theta)}{\mu_0(\theta)} - \frac{\mu(\theta')}{\mu_0(\theta')}\right)(V(\mu) - x(\mu))dp(\mu) = 0. \tag{5}$$

Obedience: for any $\mu \in \boldsymbol{supp}\{p\}$,

$$V(\mu) \in \mathbb{V}(\mu). \tag{6}$$

Bayesian Plausibility:

$$p \in BP(\mu_0). \tag{7}$$

Proof. This proof is relegated to Appendix B. □

By identifying simpler PBEs that yield the same equilibrium outcome, Proposition 2 reduces the equilibrium constraints in Equation 4 to three types of constraints. Furthermore, when we attempt to calculate $\mathcal{V}^*(\mu_0)$, we can further streamline the obedience constraints.

Corollary 1. *$\mathcal{V}^*(\mu_0)$ can be calculated by following optimization problem.*

$$\begin{aligned}
&\sup_{p\in BP(\mu_0),x} \int_\mu (V(\mu)-x(\mu))dp(\mu) \\
&\quad s.t. V(\mu)=\max \mathbb{V}(\mu) \\
&\quad \int_\mu (\frac{\mu(\theta)}{\mu_0(\theta)}-\frac{\mu(\theta')}{\mu_0(\theta')})(V(\mu)-x(\mu))dp(\mu)=0, \forall \theta,\theta'\in\Theta.
\end{aligned} \tag{8}$$

Proof. This proof is relegated to Appendix B. □

Henceforth, unless specified otherwise, let $V(\mu)$ denote $\max \mathbb{V}(\mu)$.

3 Binary-Type Illustration

In this section, we use a binary-type example to elucidate the main insights of the value of MDMB and its comparison to the value of MD which is denoted by $\mathcal{V}^*_{MD}(\mu_0)$ for prior μ_0.

We consider a salesman problem between a consumer (Receiver) and a salesman (Sender). The consumer faces a binary choice of whether to purchase a product, whose quality is either high (θ^H) or low (θ^L). The salesman has private information about the true quality of the product, while the consumer has a prior belief that the product is high-quality with probability $0<\mu_0<\frac{1}{2}$. We assume that the market price of the product is fixed at 5. The consumer's payoff depends on the quality of the product: if she purchases a high-quality product, she receives a feedback of 10; if she purchases a low-quality product, she receives a feedback of 0. The salesman's payoff is determined by the consumer's decision: he receives a payoff of 1 from the producer as his commission for selling the product and receives nothing otherwise.

In this example, we define the Sender's *interim signaling payoff* of type θ under the signaling scheme π as

$$V_\pi(\theta) \triangleq \sum_{\mu\in\mathbf{supp}\{\pi(\theta)\}} \pi(\mu|\theta)V(\mu),$$

or equivalently, suppose π induce distribution of Receiver's posteriors p, $V_p(\theta)\triangleq\int_\mu \frac{\mu(\theta)}{\mu_0(\theta)}V(\mu)dp(\mu)$. Suppose two types θ,θ' of Sender have different interim payoffs, i.e., $V_\pi(\theta)\neq V_\pi(\theta')$. Then, the type of Sender with the lower interim payoff has an incentive to deviate and input the other type. This deviation undermines the trust between Sender and Receiver and results in a loss of benefits from information disclosure of Sender. One way to mitigate this problem is to use burning money, which is a voluntary sacrifice of some benefits by Sender to gain more credibility and can be easily verified by Receiver.

In order to obtain the value of MDMB, we conduct a two-step analysis. The first step is to identify the upper bound of the salesman's expected payoff. The second step is to construct a sequence of MDMBs such that the salesman's payoff in an equilibrium under these mechanisms approaches the upper bound.

The upper bound. We begin by deriving the upper bound of the value of MDMB. Consider a mechanism (π, x) that satisfies the incentive-compatible constraint. In this mechanism, the salesman's payoff is $V_\pi(\theta^H) - \sum_{\mu \in \mathbf{supp}\{\pi(\theta^H)\}} \pi(\mu|\theta^H)x(\mu)$, which is equal to $V_\pi(\theta^L) - \sum_{\mu \in \mathbf{supp}\{\pi(\theta^L)\}} \pi(\mu|\theta^L)x(\mu)$. Since $x(\mu) \geq 0$, it follows that

$$\mathcal{V}^*(\mu_0) \leq \max_\pi\{\min\{V_\pi(\theta^H), V_\pi(\theta^L)\}\} \leq \min\{\max_\pi V_\pi(\theta^H), \max_\pi V_\pi(\theta^L)\}.$$

According to the method of Kamenica and Gentzkow [2011], we can geometrically characterize $\max_\pi V_\pi(\theta)$. To this end, we define type θ's *share of ex-post payoff* $V(\mu)$ given the posterior belief μ as

$$\hat{V}_\theta(\mu) \triangleq \frac{\mu(\theta)}{\mu_0(\theta)} V(\mu),$$

which implies that each type in the support of μ receives a share $\frac{\mu(\theta)}{\mu_0(\theta)}$ of the ex-post payoff $V(\mu)$ given the posterior belief μ. If we denote the concave envelope of function f by $cav(f)$, we can succinctly write that $\max_\pi V_\pi(\theta) = \max_p \sum_\mu p(\mu)\hat{V}(\mu) = cav(\hat{V}_\theta)(\mu_0)$. The geometric illustrations of $\max_\pi V_\pi(\theta^H), \max_\pi V_\pi(\theta^L)$ are given in Figure 2. Hence, we obtain $\mathcal{V}^*(\mu_0) \leq cav(\hat{V}_{\theta^L})(\mu_0) = \frac{\mu_0}{1-\mu_0}$.

Construction of MDMB. We construct a canonical MDMB that attains the upper bound. The signaling scheme depicted in Figure 2 is characterized by $M = \{\frac{1}{2}, 0\}$ and $\pi(\frac{1}{2}|\theta^H) = 1, \pi(\frac{1}{2}|\theta^L) = \frac{\mu_0}{1-\mu_0}, \pi(0|\theta^L) = \frac{1-2\mu_0}{1-\mu_0}$. Subsequently, we introduce a money burning message "1". For any $\delta > 0$, the modified signaling scheme is defined as $M^* = \{1, \frac{1}{2}, 0\}$, $\pi^*(\frac{1}{2}|\theta^H) = 1-\delta, \pi^*(1|\theta^H) = \delta$ and $\pi^*(\frac{1}{2}|\theta^L) = (1-\delta)\pi(\frac{1}{2}|\theta^L), \pi^*(0|\theta^L) = 1 - \pi^*(\frac{1}{2}|\theta^L)$. The associated money-burning scheme is $x(1) = \frac{1-(2-\delta)\mu_0}{\delta(1-\mu_0)}, x(\frac{1}{2}) = t(0) = 0$. This MDMB is canonical and incentive-compatible for any $\delta > 0$. The salesman's expected payoff is $\frac{\mu_0}{1-\mu_0}(1-\delta)$. As $\delta \to 0^+$, the salesman's payoff converges to the upper

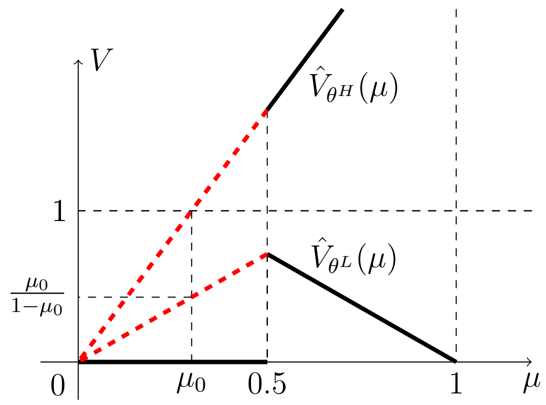

Figure 2: $\hat{V}_{\theta^H}(\mu)$ and $\hat{V}_{\theta^L}(\mu)$.

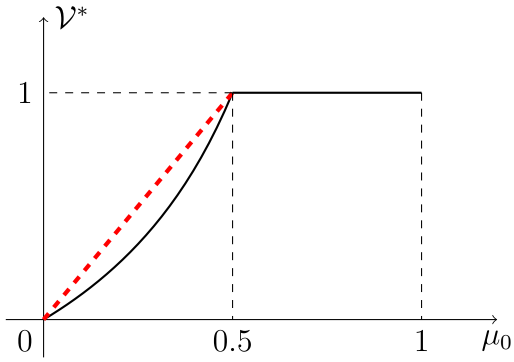

Figure 3: $\mathcal{V}^*(\mu)$.

bound of the value of MDMB. We can illustrate the value of MDMB $\mathcal{V}^*(\mu_0)$ for different priors μ_0 in Figure 3.

The comparison to the value of MD. So far, we have shown that the value of MDMB $\mathcal{V}^*(\mu_0) = \frac{\mu_0}{1-\mu_0}$ for $\mu_0 < 0.5$. On the other hand, Salamanca [2021], Corrao and Dai [2023] explore the value of MD, through which we can derive that $\mathcal{V}^*_{MD}(\mu_0) = \inf_{\lambda \in \mathbb{R}} cav(\lambda \hat{V}_{\theta^H} + (1-\lambda)\hat{V}_{\theta^L})(\mu_0) = 0$ for $\mu_0 < 0.5$. Therefore, in this example, classical MD is not credible at all (similar to cheap talk). However, the money-burning tactic provides credibility and enhances the Sender's communication efficiency in all cases where $0 < \mu_0 < 0.5$.

4 The Value of MDMB

Given a canonical MDMB (π, x), we define the *interim signaling payoff* of type θ under signaling scheme π as follows,

$$V_\pi(\theta) \triangleq \sum_{\mu \in \mathbf{supp}\{\pi(\theta)\}} \pi(\mu|\theta) V(\mu). \tag{9}$$

Moreover, we define the type θ's share of ex-post payoff $V(\mu)$ given the posterior μ as $\hat{V}_\theta(\mu) = \frac{\mu(\theta)}{\mu_0(\theta)} V(\mu)$.[7] Based on this adjusted ex-post payoff, we introduce the Sender's

[7]This follows from the fact that $\mathbb{E}_{\theta \sim \mu_0}\{\hat{V}_\theta(\mu)\} = V(\mu)$. The notation $\hat{V}_\theta$ is also known as truth-adjust welfare function introduced by Doval and Smolin [2024].

subjective payoff function under the posterior μ and the subjective prior $\lambda \in \Delta(\Theta)$ as

$$\hat{V}_\lambda(\mu) \triangleq \mathbb{E}_{\theta\sim\lambda}\{\hat{V}_\theta(\mu)\} = \sum_{\theta\in\Theta} \lambda(\theta)\frac{\mu(\theta)}{\mu_0(\theta)}V(\mu).^{8} \tag{10}$$

We denote the concave envelope of function f by $cav(f)$. Now, we can derive the value of MDMB.

Proposition 3. $\mathcal{V}^*(\mu_0) = \max_\pi \min_{\theta\in\Theta} V_\pi(\theta) = \min_{\lambda\in\Delta(\Theta)} cav(\hat{V}_\lambda)(\mu_0)$.

Proof. This proof is relegated to Appendix B. □

We have elucidated the underlying intuition of the proof of Proposition 3 in Section 3. Formally, our proof of Proposition 3 is structured into three distinct components. Initially, we leverage the non-negativity of money burning to establish the upper bound for $\mathcal{V}^*(\mu_0)$. Subsequently, for any signaling scheme π, we present a construction of the MDMB, encapsulated in Proposition 11, demonstrating that Sender can secure at least the minimum interim signaling payoff associated with π. Ultimately, the min-max and max-min equality is demonstrated through the application of Sion's minimax theorem.

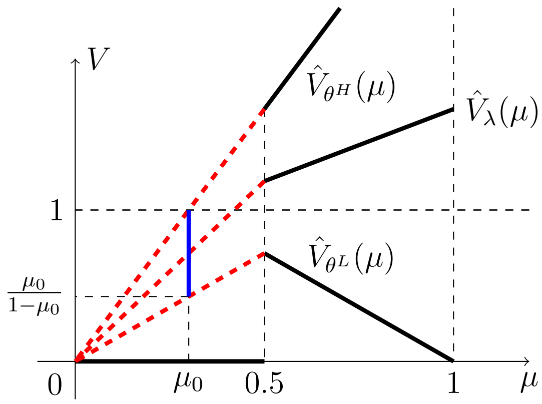

Figure 4: The geometric interpretation.

According to Proposition 3, we can geometrically characterize the value of MDMB and explain how the money burning mechanism reshapes the value function. As depicted in Figure 4, Proposition 3 characterizes the value of the MDMB for the example in Section 3. We first plot the ex-post payoff shares of type θ^H and type θ^L, denoted by $\hat{V}_{\theta^H}(\mu)$ and $\hat{V}_{\theta^L}(\mu)$, respectively, which are both exogenously given. Then we compute $\hat{V}_\lambda$, which is the convex combination or the expected value of $\hat{V}_{\theta^H}$ and $\hat{V}_{\theta^L}$ under λ. The value of the

[8] Note that when $\lambda = \mu_0$ the subjective payoff function under the posterior μ becomes $V(\mu)$.

MDMB is the minimum of $cav(\hat{V}_\lambda)(\mu_0)$ over all possible λ, as shown by the blue line in Figure 4. Therefore, the money burning mechanism transforms the belief-payoff function from $V(\mu)$ to $\hat{V}_{\theta^L}(\mu)$ in this example.

4.1 The Optimal MDMB

According to Proposition 3, the key of constructing the optimal MDMB lies in determining a signaling scheme that maximizes the minimum interim signaling payoff, which we refer to as *the optimal signaling scheme*. Once we have this optimal signaling scheme, we can apply Proposition 11 to construct the optimal MDMB. Therefore, in this section, our calculate the optimal signaling scheme according to the min-max and max-min equivalence.

We first explain the intuition for the optimal signaling scheme design. We notice that the min-mas and max-min equality can be modeled as the Nash equilibrium of a zero-sum game between Sender and Nature. Sender designs the signaling scheme $p \in BP(\mu_0)$ while Nature sets Sender's subjective prior $\lambda \in \Delta(\Theta)$. Sender aims to maximize the following payoff function in the zero-sum game while Nash pursues to minimize it.

$$\mathcal{L}(\lambda, p) \triangleq \int_\mu \hat{V}_\lambda(\mu) dp(\mu). \tag{11}$$

The following proposition provides the indifferent condition for Nature's optimal strategy at the Nash equilibrium in the zero-sum game.

Proposition 4. *A subjective prior $\lambda^* \in \Delta(\Theta)$ is the worst Sender's subjective prior if and only if there exists $p^* \in BP(\mu_0)$ such that $\mathcal{L}(\lambda^*, p^*) = cav(\hat{V}_{\lambda^*})(\mu_0)$, and for any $\theta \in \boldsymbol{supp}(\lambda^*)$, $\mathcal{L}(\lambda^*, p^*) = \mathcal{L}(\mu_\theta, p^*)$, and for any $\theta \notin \boldsymbol{supp}(\lambda^*)$, $\mathcal{L}(\lambda^*, p^*) \leq \mathcal{L}(\mu_\theta, p^*)$.*[9]

We also get the condition for Sender's optimal strategy at the Nash equilibrium and explain it in the following proposition.

Proposition 5. *A signaling scheme $p^* \in BP(\mu_0)$ is optimal if and only if there exists $\lambda^* \in \Delta(\Theta)$ such that $\mathcal{L}(\lambda^*, p^*) = cav(\hat{V}_{\lambda^*})(\mu_0)$, and for any $\theta \in \boldsymbol{supp}(\lambda^*)$, $\mathcal{L}(\lambda^*, p^*) = \mathcal{L}(\mu_\theta, p^*)$, and for any $\theta \notin \boldsymbol{supp}(\lambda^*)$, $\mathcal{L}(\lambda^*, p^*) \leq \mathcal{L}(\mu_\theta, p^*)$.*

Proof. Those proofs are relegated to Appendix B. □

[9] μ_θ is the distribution in $\Delta(\Theta)$ with a singleton support $\{\theta\}$.

These two propositions jointly characterize the worst subjective prior and the optimal signaling scheme. The condition, "$\mathcal{L}(\lambda^*, p^*) = cav(\hat{V}_{\lambda^*})(\mu_0)$, and for any $\theta \in \mathbf{supp}(\lambda^*)$, $\mathcal{L}(\lambda^*, p^*) = \mathcal{L}(\mu_\theta, p^*)$, and for any $\theta \notin \mathbf{supp}(\lambda^*)$, $\mathcal{L}(\lambda^*, p^*) \leq \mathcal{L}(\mu_\theta, p^*)$", implies that λ^* and p^* form a Nash equilibrium in the "zero-sum game", where λ^* is the best response to p^* and vice versa.

We can apply the characterization of the worst subjective prior to narrow the possible range of the set $\arg\min_{\lambda \in \Delta(\Theta)} cav(\hat{V}_\lambda)(\mu_0)$ when there are only two possible types of the Sender.

Proposition 6. *Suppose Sender has a binary type set $\Theta = \{\theta_1, \theta_2\}$. Then, for any prior distribution $\mu_0 \in \Delta(\Theta)$, we have $\mathcal{V}^*(\mu_0) = \min\{cav(\hat{V}_{\theta_1})(\mu_0), cav(\hat{V}_{\theta_2})(\mu_0)\}$.*

Proof. This proof is relegated to Appendix B. □

In Example 1, however, we show that the worst subjective prior is not necessarily an extreme point of $\Delta(\Theta)$ in general.

5 Money Burning Improves Mediated Communication

This section primarily discusses when money burning tactic can strictly improve the Sender's efficiency of MD. Let $\mathcal{V}^*_{MD}(\mu_0)$ denote the Sender's maximum equilibrium payoff under MD without money burning.

In order to introduce our approach to compare $\mathcal{V}^*_{MD}(\mu_0)$ and $\mathcal{V}^*(\mu_0)$, we need a topological generic[10] property of the Receiver's payoff function set. When the Receiver's payoff function exhibits this generic property, our communication protocol can improve the Sender's payoff under almost all Receiver's beliefs where commitment is valuable.

Definition 3. *We call the setting of A, u and Θ generic if, for any belief μ and any $a \in RO(\mu)$, there exists μ' such that $\boldsymbol{supp}\{\mu'\} = \boldsymbol{supp}\{\mu\}$ and $RO(\mu') = \{a\}$.*[11]

[10]The Lebesgue measure of the set of the Receiver's payoff that fails to satisfy this condition is zero throughout the entire payoff function space.

[11]This condition is also present in Lipnowski et al. [2024], where they employ it as a generic sufficient criterion for the uniqueness of the Sender's payoff under perfect Bayesian equilibrium in the Bayesian persuasion game.

Let $\mathcal{V}^*_{CT}(\mu_0)$ denote the Sender's maximum payoff under cheap talk equilibrium given the prior μ_0. According to Lipnowski and Ravid [2020], $\mathcal{V}^*_{CT}(\mu_0) = qcav(V)(\mu_0)$ where $qcav(f)$ is the quasi-concave envelope of function f. Then, we have following theorem to answer the question when money burning is strictly valuable to MD.

Theorem 1. *Under generic settings Definition 3, for almost all prior beliefs $\mu_0 \in \Delta(\Theta)$, either $\mathcal{V}^*_{CT}(\mu_0) = \max_{\mu\in\Delta(\Theta)} V(\mu)$ or $\mathcal{V}^*_{MD}(\mu_0) < \mathcal{V}^*(\mu_0)$.*

Clearly if $\mathcal{V}^*_{CT}(\mu_0) = \max_{\mu\in\Delta(\Theta)} V(\mu)$, commitment holds no value. Therefore, Theorem 1 demonstrates that in cases where commitment is valuable, the money-burning tactic generically leads to a strict improvement in MD when Sender has a state-independent preference.

The straightforward approach of proving Theorem 1 does not work. According to Salamanca [2021], Corrao and Dai [2023], we have the result that $\mathcal{V}^*_{MD}(\mu_0) = \min_{\lambda\in aff(\Theta)} cav(\hat{V}_\lambda)(\mu_0)$ where, for $\lambda \in aff(\Theta)$,

$$\hat{V}_\lambda(\mu_0) = \max\{\sum_{\theta\in\Theta} \lambda(\theta)\frac{\mu(\theta)}{\mu_0(\theta)} \max \mathbb{V}(\mu), \sum_{\theta\in\Theta} \lambda(\theta)\frac{\mu(\theta)}{\mu_0(\theta)} \min \mathbb{V}(\mu)\}.$$

Although we have clear characterizations of $\mathcal{V}^*(\mu_0)$ and $\mathcal{V}^*_{MD}(\mu_0)$, it remains challenging to directly ascertain when $\mathcal{V}^*(\mu_0)$ is strictly greater than $\mathcal{V}^*_{MD}(\mu_0)$ due to the lack of a clear characterization of the set $\arg\min_{\lambda\in aff(\Theta)} cav(\hat{V}_\lambda)(\mu_0)$. As a result, it is difficult to compare the sets $\arg\min_{\lambda\in aff(\Theta)} cav(\hat{V}_\lambda)(\mu_0)$ and $\arg\min_{\lambda\in\Delta(\Theta)} cav(\hat{V}_\lambda)(\mu_0)$ to determine whether the money-burning tactic improves efficiency. Moreover, for this reason, obtaining a quantifiable result regarding the extent to which the money-burning tactic enhances efficiency is also challenging.

Therefore, we adopt an alternative approach to prove Theorem 1. The proof consists of two main steps. The first step is based on an observation that if the value of MDMB is equal to the value of MD, then the value of cheap talk will also be equal to this value. The second step is a constructive proof demonstrating that, generically, $\mathcal{V}^*_{CT}(\mu_0) < \mathcal{V}^*(\mu_0)$.

Proposition 7. *Under generic settings Definition 3, if $\mathcal{V}^*(\mu_0) = \mathcal{V}^*_{MD}(\mu_0)$ then $\mathcal{V}^*_{CT}(\mu_0) = \mathcal{V}^*_{MD}(\mu_0) = \mathcal{V}^*(\mu_0)$.*

Proof. This proof is relegated to Appendix B. □

Technically, we provide an example in Appendix A. Example 3 shows that when the generic condition is not satisfied, it is possible that $\mathcal{V}^*_{CT}(\mu_0) < \mathcal{V}^*_{MD}(\mu_0) = \mathcal{V}^*(\mu_0)$, which demonstrates the necessity of generic settings of Proposition 7.

By Proposition 7, we observe that if $\mathcal{V}^*(\mu_0) > \mathcal{V}^*_{CT}(\mu_0)$, money burning tactic can strictly improve MD. Hence, based on this observation, we only need to find a sufficient condition under which $\mathcal{V}^*(\mu)$ is strictly larger than $\mathcal{V}^*_{CT}(\mu_0)$. The following proposition provides such a sufficient condition.

Proposition 8. *Under generic settings Definition 3, if $qcav(V)(\mu_0) \neq cav(V)(\mu_0)$ and there is a sufficiently small $\varepsilon > 0$ such that $qcav(V)(\mu_0 + \varepsilon(\mu - \mu_0)) = qcav(V)(\mu_0)$ for all $\mu \in \Delta(\Theta)$, it follows that $\mathcal{V}^*_{CT}(\mu_0) < \mathcal{V}^*(\mu_0)$.*

Proof. This proof is relegated to Appendix B. □

We can use Proposition 8 to identify when there is a positive improved value of mediator. The condition $qcav(\mathbb{V})(\mu_0) \neq cav(\mathbb{V})(\mu_0)$ rules out the possibility that commitment has no value. The condition that there exists a sufficiently small $\varepsilon > 0$ such that $qcav(\mathbb{V})(\mu_0 + \varepsilon(\mu - \mu_0)) = qcav(\mathbb{V})(\mu_0)$ for all $\mu \in \Delta(\Theta)$ implies that μ_0 is an interior point of the quasi-concavification distribution of the posterior of $\mathbb{V}(\mu)$ at point μ_0.

So far, we can prove Theorem 1.

Proof of Theorem 1. Remark that, under the condition that A is a finite set, $\max \mathbb{V}(\mu)$ is a piecewise constant function, and thus $qcav(\mathbb{V})(\mu)$ is also a piecewise constant function with finite pieces. Hence, following a similar argument as Corollary 2 in Lipnowski and Ravid [2020], we can deduce that *almost all* beliefs μ_0 are either interior points of some quasi-concavification distribution of posteriors or $\mathcal{V}^*_{CT}(\mu_0) = \max_{\mu \in \Delta(\Theta)} V(\mu)$. By Proposition 8, for the former μ_0, we have that $\mathcal{V}^*_{CT}(\mu_0) < \mathcal{V}^*(\mu_0)$. Therefore, by Proposition 7 we have that $\mathcal{V}^*(\mu_0) > \mathcal{V}^*_{MD}(\mu_0)$. Consequently, Theorem 1 holds. □

6 Discussions and Extensions

6.1 Relations to Bayesian Persuasion

The section discusses the relationship between Bayesian persuasion and MDMB. Let $\mathcal{V}^*_{BP}(\mu_0)$ represent the value of Bayesian persuasion, that is $\mathcal{V}^*_{BP}(\mu_0) = cav(V)(\mu_0)$. In the first part of this section, we demonstrate that the value of MDMB can be interpreted as the value of some types of robust Bayesian persuasion. In the second part, we directly compare $\mathcal{V}^*_{BP}(\mu_0)$ with $\mathcal{V}^*(\mu_0)$, which encapsulates the value of commitment within Web 3 communities.

6.1.1 Robust Bayesian persuasion

In addition to the geometric property of $\mathcal{V}^*(\mu_0)$, Proposition 3 also yields two important implications, linking the value of MDMB to two varieties of robust Bayesian persuasion problems, which provides a theoretical micro foundation for maxmin utility.

The first equation of Proposition 3 indicates that the value of MDMB is the same as the value of Sender who has full commitment power and opts for a signaling scheme that maximize his minimum interim payoff. The model of Sender opting for such a signaling scheme is named cautious Bayesian persuasion in which Sender only focuses on his lowest possible welfare, Doval and Smolin [2021, 2024]. In cautious Bayesian persuasion setting, Sender has full commitment power but acts robustly to the type realization. The following corollary constitutes our first implication of Proposition 3.

Corollary 2. *The value of MDMB equals the payoff of Sender with full commitment power but who is cautious.*

The second equation of Proposition 3 relates the value of MDMB to the Sender's payoff in robust Bayesian persuasion with heterogeneous beliefs. To elucidate this, we call a subjective distribution $\lambda \in \Delta(\Theta)$ as the *worst Sender's subjective prior* if it minimizes $cav(\hat{V}_\lambda)(\mu_0)$ which we refer to as *the worst Sender's subjective expected payoff.* Based on the model introduced by Alonso and Câmara [2016] in which the Sender's and the Receiver's subjective priors are heterogeneous, the second equation of Proposition 3 shows that the value of the MDMB coincides with the worst Sender's subjective expected payoff in heterogeneous belief Bayesian persuasion. Hence, the following corollary is our second implication of Proposition 3.

Corollary 3. *The value of MDMB $\mathcal{V}^*(\mu_0)$ equals to the payoff of Sender under Bayesian persuasion with heterogeneous priors, in which Sender holds the worst subjective prior and Receiver has prior μ_0.*

6.1.2 The Value of Commitment in Web 3 Communities

In conventional societies, the paradigm of communication without commitment is epitomized by the cheap talk model. Consequently, $\mathcal{V}^*_{BP}(\mu_0) - \mathcal{V}^*_{CT}(\mu_0)$ quantifies the value of commitment inherent in communication within conventional societies. However, as delineated in Drakopoulos et al. [2023], the paradigm of communication in Web 3.0 communities, facilitated by Blockchain systems and smart contracts, presents a radically altered landscape. In Web 3.0 communities, users are characterized by full decentralization and a potential for high unreliability. Consequently, a viable approach to facilitating communication among these users is through the deployment of smart contracts, which serve as transparent algorithms. Sender leverage smart contracts to integrate money-burning mechanisms via subsidies and gas fees. Therefore, the communication milieu of Web 3.0 communities is not amenable to modeling as cheap talk but rather as MDMB. Hence, we denote $\mathcal{V}_{BP}(\mu_0) - \mathcal{V}^*(\mu_0)$ as the *refined value of commitment* in Web 3.0 communities.

Our first result establishes the condition under which the refined value of commitment does not exist.

Proposition 9. *If $\mathcal{V}^*(\mu_0) = \mathcal{V}^*_{BP}(\mu_0)$ then $\mathcal{V}^*_{CT}(\mu_0) = \mathcal{V}^*_{MD}(\mu_0) = \mathcal{V}^*(\mu_0) = \mathcal{V}^*_{BP}(\mu_0)$.*

Proof. This proof is relegated to Appendix B. □

Intuitively, we can derive this proposition from the result of Corrao and Dai [2023]. When $\mathcal{V}^*(\mu_0) = \mathcal{V}^*_{BP}(\mu_0)$, there is no money burning. Hence, $\mathcal{V}^*_{MD}(\mu_0) = \mathcal{V}^*_{BP}(\mu_0)$, which leads to $\mathcal{V}^*_{CT}(\mu_0) = \mathcal{V}^*_{BP}(\mu_0)$ by Corrao and Dai [2023]. By Proposition 9, we deduce that if $\mathcal{V}^*_{CT}(\mu_0) < \mathcal{V}^*_{BP}(\mu_0)$, then $\mathcal{V}^*(\mu_0) < \mathcal{V}^*_{BP}(\mu_0)$. Consequently, we promptly arrive at the subsequent corollary, which juxtaposes the value of commitment with the refined value of commitment.

Corollary 4. *There is a positive value of commitment, i.e. $\mathcal{V}^*_{CT}(\mu_0) < \mathcal{V}^*_{BP}(\mu_0)$, if and only if there is a positive refined value of commitment in the Web 3.0 community, i.e. $\mathcal{V}^*(\mu_0) < \mathcal{V}^*_{BP}(\mu_0)$.*

Technically, we provide an example in Appendix A. Example 2 shows that the relation $\mathcal{V}^*_{CT}(\mu_0) = \mathcal{V}^*_{MD}(\mu_0) = \mathcal{V}^*(\mu_0) < \mathcal{V}^*_{BP}(\mu_0)$ can hold even under generic settings, which refers to the case that the refined value of commitment is positive and the same as the value of commitment.

Corollary 4 has an important implication that if there is the value of commitment in conventional societies, commitment is still valuable in Web 3.0 communities. In addition, according to Corollary 2 in Lipnowski and Ravid [2020], the refined value of commitment is strictly positive under almost all prior beliefs as long as the value of commitment is strictly positive. Furthermore, building on Corollary 2 and Corollary 3, we obtain two characterizations of the refined value of commitment. Corollary 2 indicates that the refined value of commitment is given by the difference between BP and cautious BP. This difference can be assessed by comparing the expected payoff of the Sender and the lowest interim payoff of the Sender under full commitment. Corollary 3 implies that the refined value of commitment is also given by the difference between BP and heterogeneous BP under the Sender's worst subjective prior. This difference can be evaluated by considering the Receiver's prior μ_0 and the Sender's worst subjective prior λ^*.

6.2 Relations to Money Burning in Cheap Talk

Numerous studies have explored the role of money burning in cheap talk equilibria Austen-Smith and Banks [2000], Kartik [2007], Karamychev and Visser [2017], where it can be viewed as a mechanism for designing a costly signaling game. These studies reveal that while money burning can enhance the precision of a cheap talk equilibrium, it only makes cheap talk influential[12] if some cheap talk equilibrium themselves, without money burning, are already influential.

Under the assumption of transparent motives, these results apply directly, demonstrating that the Sender cannot leverage money burning to improve communication efficiency in a cheap talk equilibrium. In the binary example provided in Section 3, we observe that, under a prior of $\mu_0 < 0.5$, the Sender cannot achieve a higher overall payoff than the no-communication payoff by using protocols such as cheap talk, cheap talk with money burning, or mediated communication without money burning. More specifically, these communication protocols fail to produce an influential equilibrium. However, in our construction in Section 3, money burning does indeed expand the credibility boundary of

[12] An equilibrium is considered influential if there are at least two of the Receiver's actions taken along the equilibrium path with the same amount of money burning.

mediated communication, resulting in an influential equilibrium. This distinction highlights the differing roles of money burning in cheap talk versus mediated communication. As noted by Theorem 1 in Austen-Smith and Banks [2000], money burning can enhance the precision of cheap talk communication but does not improve its credibility. In contrast, Proposition 7 demonstrates that money burning can indeed bolster the credibility of mediated communication.

Beyond the assumption of transparent motives, a study on the implementation problem Liu and Wu [2024] demonstrates a significant difference in the implementability conditions between implementing a cheap talk equilibrium with money burning and a mediated communication equilibrium with money burning. In addition, in Example 4, we show that money burning cannot make mediated communication credible in any case, implying that money burning cannot enhance mediated communication and the necessity of the transparent motives assumption.

6.3 Bounded Credit for Money Burning

This section addresses the scenario where the Sender's money burning is limited by a budget constraint C. Formally, the MDMB with a budget constraint C for money burning is characterized by an input set M, an output set S, and a mapping $\phi : M \to \Delta(S\times[0, C])$. The Sender-Receiver game remains identical to that explained in Section 2.1. For any given value of C, we use $\mathcal{V}^*_C(\mu_0)$ to denote the value of MDMB with a budget constraint C for money burning.

To analyze the optimal MDMB with budget constraint, we define a new concept of Sender's *generalized ex-post subjective payoff function* for a posterior $\mu \in \Delta(\Theta)$, denoted by $\hat{V}_{\lambda,C}(\mu)$, as follow

$$\hat{V}_{\lambda,C}(\mu) \triangleq \max\{\sum_{\theta\in\Theta} \lambda(\theta)\frac{\mu(\theta)}{\mu_0(\theta)} \max \mathbb{V}(\mu), \sum_{\theta\in\Theta} \lambda(\theta)\frac{\mu(\theta)}{\mu_0(\theta)}(\min \mathbb{V}(\mu) - C)\}. \tag{12}$$

Here, $\lambda : \Theta \to \mathbb{R}$ is the Lagragian multiplier of incentive-compatible constraints and satisfies $\sum_\theta \lambda(\theta) = 1$. Thus, $\hat{V}_{\lambda,C}(\mu)$ is an affine combination of the adjusted share of Sender's ex-post payoff similar to Equation 10. Then, we use $\hat{V}_{\lambda,C}(\mu)$ to define the *generalized ex-ante subjective expected payoff function* associated with the signal scheme

$p \in BP(\mu_0)$ as below.

$$\mathcal{L}_C(\lambda, p) \triangleq \int_\mu \hat{V}_{\lambda,C}(\mu) dp(\mu). \tag{13}$$

Then, we have the following theorem characterizing the value of MDMB with budget constraints.

Proposition 10.

$$\mathcal{V}_C^*(\mu_0) = \max_{p \in BP(\mu_0)} \min_{\lambda \in \mathit{aff}(\Theta)} \mathcal{L}_C(\lambda, p) = \min_{\lambda \in \mathit{aff}(\Theta)} \max_{p \in BP(\mu_0)} \mathcal{L}_C(\lambda, p).$$

Moreover, $\mathcal{V}_C^*(\mu_0) = \min_{\lambda \in \mathit{aff}(\Theta)} cav(\hat{V}_{\lambda,C})(\mu_0)$.

Proof. This proof is relegated to Appendix B. □

To establish Proposition 10, we initially generalize the revelation principle (Proposition 1) to demonstrate that computing $\mathcal{V}_C^*(\mu_0)$ in the bounded credit environment necessitates the incorporation of an additional constraint, $x(\mu) \leq C$ for all μ, into the Sender's optimization problem (Equation 8). Subsequently, we use a two-step optimization approach, commencing with the determination of a money burning scheme that maximizes the Sender's payoff for any given signaling scheme. This process facilitates the identification of the optimal choice of obedience constraints from the Sender's perspective. This stage culminates in a max-min characterization. The subsequent step involves maximizing the Sender's payoff across signaling schemes, which is achieved through the application of Sion's minimax theorem.

The structure of $\hat{V}_{\lambda,C}$ shows that we can classify the messages of an optimal MDMB into two groups. One is the persuasion group that aims at persuasion and is costless, which corresponds to the case where $\sum_\theta \lambda^*(\theta)\frac{\mu(\theta)}{\mu_0(\theta)} \geq 0$. The other is the commitment-gaining group that aims at maximizing the credibility that Sender can gain by sacrificing his payoff, which corresponds to the case where $\sum_\theta \lambda^*(\theta)\frac{\mu(\theta)}{\mu_0(\theta)} < 0$. When the message is selected from the persuasion group, Sender desires Receiver to take the best response that favors him; Otherwise, Sender maximizes the amount of money burning and desires Receiver to take the action that mostly disfavors him.

At the end of this section, we apply Proposition 10 to compare the MDMBs with different bounded credit for money burning, Bayesian persuasion Kamenica and Gentzkow [2011], mediated communication Salamanca [2021], and cheap talk Lipnowski and Ravid

[2020] in the illustrative example in Section 3.

We begin by characterizing $\mathcal{V}_C^*(\mu_0)$. By Proposition 10, we need to consider four lines for a given parameter $\lambda \in \mathbb{R}$, namely $l_1(\mu) = 0, l_2(\mu) = -C(\frac{\lambda\mu}{\mu_0} + \frac{(1-\lambda)(1-\mu)}{1-\mu_0}), l_3(\mu) = \frac{\lambda\mu}{\mu_0} + \frac{(1-\lambda)(1-\mu)}{1-\mu_0}$ and $l_4(\mu) = (1-C)(\frac{\lambda\mu}{\mu_0} + \frac{(1-\lambda)(1-\mu)}{1-\mu_0})$. Correspondingly, we have that for $\mu \in [0,1]$,

$$\hat{V}_{\lambda,C}(\mu) = \begin{cases} \max\{l_1(\mu), l_2(\mu)\} & \mu < \frac{1}{2} \\ \max\{l_2(\mu), l_3(\mu)\} & \mu = \frac{1}{2} \\ \max\{l_3(\mu), l_4(\mu)\} & \mu > \frac{1}{2} \end{cases}.$$

Since $\hat{V}_{\lambda,C}$ is convex and upper semi-continuous on $\mu \in [0, \frac{1}{2})$ and $\mu \in (\frac{1}{2}, 1]$, to compute $cav(\hat{V}_{\lambda,C})(\mu_0)$ we only need to evaluate $\hat{V}_{\lambda,C}(0) = \max\{0, -C\frac{1-\lambda}{1-\mu_0}\}$, $\hat{V}_{\lambda,C}(\frac{1}{2}) = \max\{\frac{-C}{2}(\frac{\lambda}{\mu_0} + \frac{1-\lambda}{1-\mu_0}), \frac{1}{2}(\frac{\lambda}{\mu_0} + \frac{1-\lambda}{1-\mu_0})\}$ and $\hat{V}_{\lambda,C}(1) = \max\{\frac{\lambda}{\mu_0}, (1-C)\frac{\lambda}{\mu_0}\}$. Assuming $\mu_0 < \frac{1}{2}$, we can only partition μ_0 into $0, \frac{1}{2}$ or $0, 1$. Since $\hat{V}_{\lambda,C}(0), \hat{V}_{\lambda,C}(\frac{1}{2}), \hat{V}_{\lambda,C}(1)$ are all decreasing in λ for $\lambda \geq 0$, to find the minimum concavification value, we only need to consider the case of $\lambda \leq 0$. We then divide this case into two subcases: $\lambda \in [-\frac{\mu_0}{1-2\mu_0}, 0]$ and $\lambda \in (-\infty, -\frac{\mu_0}{1-2\mu_0})$. We can solve for the result and obtain that for $\mu_0 < \frac{1}{2}$,

$$\mathcal{V}_C^*(\mu_0) = \begin{cases} 0 & C \leq 1 \\ \frac{(C-1)\mu_0}{C(1-\mu_0)-\mu_0} & C > 1 \end{cases}.$$

For $\mu_0 \geq \frac{1}{2}$, $\mathcal{V}_C^*(\mu_0) = 1$.

If we apply different communication protocols to this example, we can obtain the optimal payoff corresponding to the prior μ_0, which is shown in Figure 5. The red line is the concave envelope of $\max \mathbb{V}(\mu)$, which is the result of Bayesian persuasion. The black line is the result of the MDMB with bound $C = +\infty$, i.e. $\mathcal{V}^*(\mu_0)$. The blue line is the result of $\mathcal{V}_2^*(\mu_0)$. Finally, we can see that regardless of what we use among the MDMB with bounded credit $C \leq 1$, cheap talk or classical mediated communication, we can only get the results as the green line, which cannot benefit from those protocols.

7 Concluding Remarks

In this paper, we introduce and investigate a novel communication protocol called mediated communication with money-burning mechanism (MDMB). In our communication protocol, Sender not only employs a trusted mediator to convey the message but

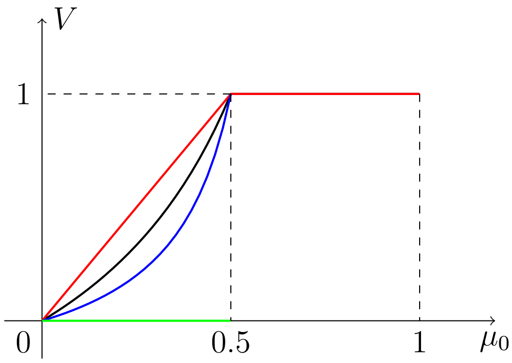

Figure 5: This figure compares the optimal payoffs of different protocols. The red line represents the payoff achieved by Bayesian persuasion. The black line shows $\mathcal{V}^*(\mu_0)$. The blue line indicates $\mathcal{V}_2^*(\mu_0)$. The green line corresponds to $\mathcal{V}_C^*(\mu_0)$ for $C \leq 1$, cheap talk, or classical mediated communication.

also voluntarily gives up some benefits to gain more commitment power. By generalizing the revelation principle of mechanism design with limited commitment, we characterize the communication efficiency of our communication protocol under the transparent motives assumption. We demonstrate that the value of MDMB aligns with that of Cautious Bayesian persuasion, which is equivalent to the concavification value under the worst Sender's subject prior. Moreover, we extend our analyses to scenarios where Sender possesses bounded credits for money-burning.

The primary finding of this paper is that, under the transparent motives assumption, money burning can strictly improve the Sender's communication efficiency in mediated communication in almost all cases where commitment holds value. Our results indicate that in balancing between money burning and enhanced commitment power, Sender generically benefits from engaging in money burning. This outcome addresses the questions of when money burning is valuable and why it is necessary in mediated communication, demonstrating a distinct role for money burning in boosting the credibility of mediated communication compared to its function in cheap talk. Additionally, this sets a new benchmark for the communication efficiency of an unreliable Sender, as money burning represents an action the Sender can undertake independently to enhance credibility.

Furthermore, the MDMB framework applies directly to Web 3.0 communities, where our findings highlight a refined concept of commitment value. We demonstrate that this refined commitment is valuable if and only if commitment is valuable without the cryptographic infrastructure. Additionally, communication efficiency in Web 3.0 communities is generically higher than in conventional settings. This finding highlights the importance of algorithmic and cryptographic technologies.

It is important to note, however, that these conclusions rely on the significant assumption of state-independent preferences for Sender. In a broader context, the role of money burning warrants further investigation. Additionally, as discussed in Section 5, a quantifiable result of the improvements from money burning remains an open problem.

References

Ricardo Alonso and Odilon Câmara. Bayesian persuasion with heterogeneous priors. *Journal of Economic Theory*, 165:672–706, 2016.

Ivan Arandjelović. An extension of the sion's minimax theorem. *Zbornik radova Filozofskog fakulteta u Nišu. Serija Matematik*, pages 1–3, 1992.

David Austen-Smith and Jeffrey S Banks. Cheap talk and burned money. *Journal of Economic Theory*, 91(1):1–16, 2000.

Dirk Bergemann and Stephen Morris. Information design: A unified perspective. *Journal of Economic Literature*, 57(1):44–95, 2019.

Dirk Bergemann, Alessandro Bonatti, and Alex Smolin. The design and price of information. *American Economic Review*, 108(1):1–48, 2018.

Helmut Bester and Roland Strausz. Contracting with imperfect commitment and the revelation principle: the single agent case. *Econometrica*, 69(4):1077–1098, 2001.

Archishman Chakraborty and Rick Harbaugh. Persuasion by cheap talk. *American Economic Review*, 100(5):2361–2382, 2010.

Roberto Corrao. Mediation markets: The case of soft information. *Working Paper*, 2023.

Roberto Corrao and Yifan Dai. The bounds of mediated communication. *Working Paper*, 2023.

Vincent P Crawford and Joel Sobel. Strategic information transmission. *Econometrica: Journal of the Econometric Society*, pages 1431–1451, 1982.

Laura Doval and Vasiliki Skreta. Mechanism design with limited commitment. *Econometrica*, 90(4):1463–1500, 2022.

Laura Doval and Alex Smolin. Information payoffs: An interim perspective. *arXiv preprint arXiv:2109.03061*, 2021.

Laura Doval and Alex Smolin. Persuasion and welfare. *Journal of Political Economy*, 132(7):2451–2487, 2024.

Kimon Drakopoulos, Irene Lo, and Justin Mulvany. Blockchain mediated persuasion. In *Proceedings of the 24th ACM Conference on Economics and Computation*, 2023.

Shaddin Dughmi, Rad Niazadeh, Alexandros Psomas, and S Matthew Weinberg. Persuasion and incentives through the lens of duality. In *Web and Internet Economics: 15th International Conference, WINE 2019, New York, NY, USA, December 10–12, 2019, Proceedings 15*, pages 142–155. Springer, 2019.

Francoise Forges. An approach to communication equilibria. *Econometrica: Journal of the Econometric Society*, pages 1375–1385, 1986.

Guillaume R Fréchette, Alessandro Lizzeri, and Jacopo Perego. Rules and commitment in communication: An experimental analysis. *Econometrica*, 90(5):2283–2318, 2022.

Maria Goltsman, Johannes Hörner, Gregory Pavlov, and Francesco Squintani. Mediation, arbitration and negotiation. *Journal of Economic Theory*, 144(4):1397–1420, 2009.

Maxim Ivanov. Beneficial mediated communication in cheap talk. *Journal of Mathematical Economics*, 55:129–135, 2014.

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

Vladimir Karamychev and Bauke Visser. Optimal signaling with cheap talk and money burning. *International Journal of Game Theory*, 46(3):813–850, 2017.

Navin Kartik. A note on cheap talk and burned money. *Journal of Economic Theory*, 136(1):749–758, 2007.

Anton Kolotilin and Hongyi Li. Relational communication. *Theoretical Economics*, 16 (4):1391–1430, 2021.

Hao Li and Xianwen Shi. Discriminatory information disclosure. *American Economic Review*, 107(11):3363–3385, 2017.

Xiao Lin and Ce Liu. Credible persuasion. *Journal of Political Economy*, 132(7):2228–2273, 2024.

Elliot Lipnowski and Doron Ravid. Cheap talk with transparent motives. *Econometrica*, 88(4):1631–1660, 2020.

Elliot Lipnowski, Doron Ravid, and Denis Shishkin. Persuasion via weak institutions. *Journal of Political Economy*, 130(10):2705–2730, 2022.

Elliot Lipnowski, Doron Ravid, and Denis Shishkin. Perfect bayesian persuasion. *arXiv preprint arXiv:2402.06765*, 2024.

Yi Liu and Fan Wu. Implementing randomized allocation rules with outcome-contingent transfers. *Journal of Economic Theory*, 220:105878, 2024.

Daehong Min. Bayesian persuasion under partial commitment. *Economic Theory*, 72(3): 743–764, 2021.

Roger B Myerson. Optimal coordination mechanisms in generalized principal–agent problems. *Journal of Mathematical Economics*, 10(1):67–81, 1982.

Eduardo Perez-Richet and Vasiliki Skreta. Test design under falsification. *Econometrica*, 90(3):1109–1142, 2022.

Hitoshi Sadakane. Multistage information transmission with voluntary monetary transfers. *Theoretical Economics*, 18(1):267–301, 2023.

Andrés Salamanca. The value of mediated communication. *Journal of Economic Theory*, 192:105191, 2021.

Monir Shaker, Fereidoon Shams Aliee, and Reza Fotohi. Online rating system development using blockchain-based distributed ledger technology. *Wireless Networks*, 27(3): 1715–1737, 2021.

Appendix

A Omitted Examples

A.1 The Worst Subjective Prior Beyond Binary Types

Example 1. *We consider an example with three parties: a seller, a buyer, and an influencer. The seller wants to sell a zero-cost product to the buyer. The buyer's valuation of the product is v. The seller only knows that v is distributed uniformly in $\{1, 2, 3\}$. The buyer is a fan of the influencer, who wants to help the buyer reduce the price of the product by disclosing information about the buyer's type and subsidizing the seller. The influencer acts as Sender who uses our MDMB to influence the seller's action as Receiver. To fit our model, we let the type set be $\Theta = \{v_1 = 1, v_2 = 2, v_3 = 3\}$, the prior distribution be $\mu_0 = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$, and the seller's action set be $A = \{p_1 = 1, p_2 = 2, p_3 = 3\}$.*

We assume that the influencer's objective is to minimize the price of the product. If the seller charges a price p_i to the buyer, the influencer's valuation function is $v(p_i) = 4 - p_i$.

We use Proposition 3 to examine the extreme point subjective priors of $\Delta(\Theta)$ at first. We then use Proposition 4 to find the worst Sender's subjective prior and the corresponding maximum payoff of the influencer achieved by the MDMB. We also derive the optimal MDMB backward.

We consider three extreme point subjective priors $\lambda_i \in \Delta(\Theta), i = 1, 2, 3$, where $\lambda_1 = (1, 0, 0), \lambda_2 = (0, 1, 0), \lambda_3 = (0, 0, 1)$. For any $\lambda \in \Delta(\Theta)$, to find the concavification value of $\hat{V}_\lambda(\mu)$ at μ_0, we can assume without loss of generality that we only need to find the distribution of posterior $\tau \in BP(\mu_0)$ that induces different actions of the Receiver.[13] Then finding $cav(\hat{V}_\lambda)(\mu_0)$ becomes a linear programming problem. We obtain that $cav(\hat{V}_{\lambda_1})(\mu_0) = 3$, where $\{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}$ forms the support of Receiver's posterior distribution and they are realized with equal probability; $cav(\hat{V}_{\lambda_2})(\mu_0) = 3$, where $(\frac{1}{2}, \frac{1}{2}, 0), (0, 0, 1)$ forms the support of Receiver's posterior distribution and they are realized with probability $\frac{2}{3}, \frac{1}{3}$ respectively; and $cav(\hat{V}_{\lambda_3})(\mu_0) = \frac{8}{3}$. So we can conclude that in this example $\mathcal{V}^*(\mu_0) \leq \frac{8}{3}$.

However, λ_3 is not the worst subjective prior in this case, even though it minimizes

[13] This is true because if two posteriors μ_1, μ_2 in the support of τ lead to the same action of the Receiver, we can merge them as posterior $\frac{\tau(\mu_1)}{\tau(\mu_1)+\tau(\mu_2)}\mu_1 + \frac{\tau(\mu_2)}{\tau(\mu_1)+\tau(\mu_2)}\mu_2$ with probability $\tau(\mu_1) + \tau(\mu_2)$.

$cav(\hat{V}_\lambda)(\mu_0)$ among the extreme points of $\Delta(\Theta)$. Next we show that $\lambda^* = (0, \frac{1}{2}, \frac{1}{2})$ is the worst subjective prior by Proposition 4. We first calculate that $cav(\hat{V}_{\lambda^*}) = \frac{5}{2}$ and the process of concavification splits μ_0 into $(\frac{1}{2}, \frac{1}{4}, \frac{1}{4}), (0, \frac{1}{2}, \frac{1}{2})$ with probability $\frac{2}{3}, \frac{1}{3}$ respectively. We denote this distribution over posterior as τ^*. We verify that $\mathcal{L}(\lambda_2, \tau^*) = \mathcal{L}(\lambda_3, \tau^*) = \frac{5}{2} < 3 = \mathcal{L}(\lambda_1, \tau^*)$. So by Proposition 4, λ^* is the worst subjective prior and by Proposition 3 we have $\mathcal{V}^*(\mu_0) = \frac{5}{2}$. Moreover, by Proposition 5, we know that τ^* is the optimal signaling scheme. We can use Proposition 11 to construct the optimal MDMB that approaches $\mathcal{V}^*(\mu_0)$.

A.2 No Value of MDMB

Example 2. *Receiver has three possible actions a_1, a_2, a_3 and Sender has two possible types H, L. The prior belief assigns probability μ_0 to the Sender's type being H. The Sender's values for the actions are $v(a_1) = 0, v(a_2) = \frac{1}{4}, v(a_3) = 1$. We summarize the Receiver's payoffs in Table 1.*

$u(a,\theta)$	H	L
a_1	-4	1
a_2	0	0
a_3	1	-2

Table 1: Receiver's payoff matrix.

In this example, the belief-value correspondence is

$$\mathbb{V}(\mu) = \begin{cases} 1 & \mu \in (\frac{2}{3}, 1] \\ [\frac{1}{4}, 1] & \mu = \frac{2}{3} \\ \frac{1}{4} & \mu \in (\frac{1}{5}, \frac{2}{3}) \\ [0, \frac{1}{4}] & \mu = 1/5 \\ 0 & \mu < \frac{1}{5} \end{cases}.$$

We depict $\mathbb{V}(\mu)$ on Figure 6 as the black line, which corresponds to the outcome under a mediator without money burning and cheap talk, following the results of Salamanca [2021]. Based on Proposition 6, we display the result of $\mathcal{V}^*$ on Figure 6 as the blue line and $\mathcal{V}^*_{BP}$ as the red line, with the procedure omitted. We observe that, when $\mu_0 = \frac{1}{5}$, the case of $\mathcal{V}^*_{CT}(\mu_0) = \mathcal{V}^*_{MD}(\mu_0) = \mathcal{V}^*(\mu_0) < \mathcal{V}^*_{BP}(\mu_0)$ arises under the generic setting.

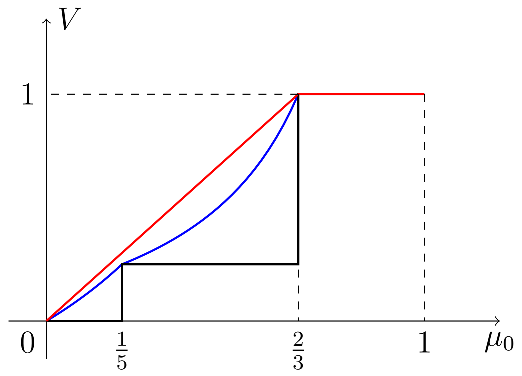

Figure 6: Results of Example 2.

A.3 The Necessity of Generic Condition

Example 3. *We present an abstract setting in this example, where we only specify the belief-value function and ensure the existence of the basic settings of A, u, v, Θ by imposing the upper-semi continuity of the belief-value function.*

We assume that there are three distinct types θ_1, θ_2 and θ_3. The maximum of belief-value correspondence is

$$V(\mu) = \begin{cases} \frac{7}{3} & \mu(\theta_1) = 1 \\ 2 & \mu(\theta_1) = 0, \mu(\theta_2) \in [0, \frac{1}{2}) \\ 3 & \mu(\theta_1) = 0, \mu(\theta_2) \in [\frac{1}{2}, \frac{3}{4}] \\ 1 & \mu(\theta_1) = 0, \mu(\theta_2) \in (\frac{3}{4}, 1] \\ 0 & \text{otherwise} \end{cases}.$$

By restricting the support to $\{\theta_2, \theta_3\}$, the value function of Example 3 coincides with Example 3 or Figure 7 in Salamanca [2021]. For $\mu_0 = (\frac{1}{2}, \frac{1}{6}, \frac{1}{3})$, since $\mathcal{V}^*_{MD}((0, \frac{1}{3}, \frac{2}{3})) = \frac{7}{3}$ as shown by Salamanca [2021], splitting μ_0 into $(1, 0, 0)$ and $(0, \frac{1}{3}, \frac{2}{3})$ yields $\mathcal{V}^*_{MD}(\mu_0) = \frac{7}{3}$. Furthermore, the interim payoff of θ_1 cannot exceed $\frac{7}{3}$, implying that $\mathcal{V}^*(\mu_0) \leq \frac{7}{3}$. Hence, we obtain $\mathcal{V}^*(\mu_0) = \mathcal{V}^*_{MD}(\mu_0) = \frac{7}{3}$. However, to find a cheap talk equilibrium with $\frac{7}{3}$ as the Sender's payoff, we need to split μ_0 into $(1, 0, 0)$ and $(0, \frac{1}{3}, \frac{2}{3})$ and keep $(1, 0, 0)$ unchanged. Since $\mathcal{V}^*_{CT}((0, \frac{1}{3}, \frac{2}{3})) = 2 < \frac{7}{3} = V((1, 0, 0))$, no cheap talk equilibrium achieves $\frac{7}{3}$ for the Sender, and thus $\mathcal{V}^*_{CT}(\mu_0) < \mathcal{V}^*_{MD}(\mu_0) = \mathcal{V}^*(\mu_0)$ for $\mu_0 = (\frac{1}{2}, \frac{1}{6}, \frac{1}{3})$.

A.4 The Necessity of Transparent-motives Assumption

Example 4. *Consider a buyer (Sender) and a seller (Receiver) in a market. The buyer's valuation of the seller's product, which is the buyer's private information, is drawn from a set $\Theta = \{1, 2\}$. The probability that the valuation v is 2 is given by μ_0.*

The seller can set a price p chosen from the set $A = \{1, 2\}$. When the seller sets a price p and the buyer's valuation is v, the payoffs are defined as follows:

- ***Buyer's payoff**: $v(p, v) = \max\{0, v - p\}$, representing the buyer's surplus.*
- ***Seller's payoff**: $u(p, v) = \mathbb{I}(v \geq p) \cdot p$, where $\mathbb{I}(v \geq p)$ is an indicator function equal to 1 if $v \geq p$ and 0 otherwise, representing the seller's revenue.*

In this example, we show that for any $\mu_0 > 0.5$, there is no MDMB that can induce an outcome where the seller sets a price of 1. Therefore, MDMB cannot improve the buyer's total payoff; however, BP can.

Because we can merge the posteriors that induce the same action of the seller, it is without loss of generality to assume that there are two possible posteriors $x_1 < x_2$ of the seller induced by an MDMB. If the MDMB is influential, then we must have $x_1 \leq 0.5 \leq x_2$.

Because we can merge the posteriors that induce the same action of seller, it is without loss of generality to assume that there are two possible posteriors $x_1 < x_2$ of the seller induced by an MDMB. If the MDMB is influential, then we must have $x_1 \leq 0.5 < \mu_0 \leq x_2$.

If the unconditional probability of the posterior x_1 is p, then we must have $px_1 + (1 - p)x_2 = \mu_0$. Suppose the expected amount of money burned by a valuation-1 buyer is $T(1)$, and the expected amount of money burned by a valuation-2 buyer is $T(2)$. Then, the incentive-compatible constraints can be written as

$$T(2) \geq T(1)$$

and

$$\frac{px_1}{px_1 + (1-p)x_2} - T(2) \geq \frac{p(1-x_1)}{p(1-x_1) + (1-p)(1-x_2)} - T(1).$$

Thus, we can deduce that

$$\frac{px_1}{px_1 + (1-p)x_2} \geq \frac{p(1-x_1)}{p(1-x_1) + (1-p)(1-x_2)},$$

which is equivalent to $x_1 \geq x_2$. A contradiction!

B Omitted Proofs

Appendix B collects all the proofs from the main body of this paper.

B.1 Omitted Proofs in Section 2.2

Proof of Proposition 1. For any MDMB (M, S, ϕ) and a corresponding PBE assessment $(\sigma, \alpha, \mu) \in \mathcal{E}[\mathcal{G}_{M,S,\phi}(\mu_0)]$, we will directly construct a canonical MDMB (π, x) and a corresponding PBE canonical assessment $(\sigma^*, \alpha^*, \mu^*) \in \mathcal{E}[\mathcal{G}_{(\pi,x)}(\mu_0)]$ such that the expected payoffs of Sender in both assessments are the same.

The canonical MDMB we constructed is as follows: for any $\mu \in \Delta(\Theta), \theta \in \Theta$,

$$\pi(\mu|\theta) = \sum_{s\in S, t\geq 0, \mu(s,t)=\mu, m\in M} \phi(s,t|m)\sigma(m|\theta), \tag{14}$$

and for any $\mu \in \Delta(\Theta)$

$$x(\mu) = \begin{cases} \frac{\sum_{s\in S,t\geq 0,\mu(s,t)=\mu,m\in M,\theta\in\Theta} \phi(s,t|m)\sigma(m|\theta)\mu_0(\theta)t}{\sum_{s\in S,t\geq 0,\mu(s,t)=\mu,m\in M,\theta\in\Theta} \phi(s,t|m)\sigma(m|\theta)\mu_0(\theta)} & \sum_{s,t,\mu(s,t)=\mu,m,\theta} \phi(s,t|m)\sigma(m|\theta)\mu_0(\theta) \neq 0 \\ 0 & \text{otherwise} \end{cases}. \tag{15}$$

Note that the above canonical MDMB is well-defined since the support of $\phi(m)$ is finite.

The canonical assessment $(\sigma^*, \alpha^*, \mu^*)$ we constructed is as follows: for all θ, $\sigma^*(\theta|\theta) = 1$, for all $\mu \in \Delta(\Theta)$, $\mu^*(\mu) = \mu$, for all $\mu \in \textbf{supp}\{\pi(\theta)\}$ for some $\theta \in \Theta$

$$\alpha^*(\mu) = \sum_{s\in S,t\geq 0,\mu(s,t)=\mu} \frac{\sum_{\theta\in\Theta,m\in M} \mu_0(\theta)\sigma(m|\theta)\phi(s,t|m)}{\sum_{s'\in S,t'\geq 0,\mu(s',t')=\mu}\sum_{\theta\in\Theta,m\in M} \mu_0(\theta)\sigma(m|\theta)\phi(s',t'|m)}\alpha(s,t).$$

and for $\mu \notin \textbf{supp}\{\pi(\theta)\}$ for any $\theta \in \Theta$ $\alpha^*(\mu)$ is any best response given posterior belief μ.

Subsequently, we will verify that the canonical assessments $(\sigma^*, \alpha^*, \mu^*) \in \mathcal{E}[\mathcal{G}_{\pi,x}(\mu_0)]$ and the payoffs of Sender in that canonical assessments and original assessments are the same.

Before the verification, we prove the following lemma.

Lemma 1. *Suppose $\mu = \mu(s,t) \in$ **supp**$\{\pi(\hat{\theta})\}$ for some $\hat{\theta}$, then*

$$\frac{\sum_{\theta\in\Theta,m\in M}\mu_0(\theta)\sigma(m|\theta)\phi(s,t|m)}{\sum_{\mu(s',t')=\mu}\sum_{\theta\in\Theta,m\in M}\mu_0(\theta)\sigma(m|\theta)\phi(s',t'|m)} = \frac{\sum_{m\in M}\sigma(m|\hat{\theta})\phi(s,t|m)}{\sum_{\mu(s',t')=\mu}\sum_{m\in M}\sigma(m|\hat{\theta})\phi(s',t'|m)}.$$

Proof of Lemma 1. By the theorem on equal ratios, it is sufficed to show that for any $\bar{\theta} \in \mathbf{supp}\{\mu\}$, we have that

$$\frac{\sum_{m\in M}\sigma(m|\hat{\theta})\phi(s,t|m)}{\sum_{\mu(s',t')=\mu}\sum_{m\in M}\sigma(m|\hat{\theta})\phi(s',t'|m)} = \frac{\sum_{m\in M}\sigma(m|\bar{\theta})\phi(s,t|m)}{\sum_{\mu(s',t')=\mu}\sum_{m\in M}\sigma(m|\bar{\theta})\phi(s',t'|m)}. \tag{16}$$

According to Bayes updating, for any $s'' \in S, t'' \geq 0$ such that $\mu(s,t) = \mu(s'',t'')$, we have that

$$\frac{\mu_0(\hat{\theta})\sum_{m\in M}\sigma(m|\hat{\theta})\phi(s,t|m)}{\mu_0(\bar{\theta})\sum_{m\in M}\sigma(m|\bar{\theta})\phi(s,t|m)} = \frac{\mu(\hat{\theta}|s,t)}{\mu(\bar{\theta}|s,t)} = \frac{\mu(\hat{\theta}|s'',t'')}{\mu(\bar{\theta}|s'',t'')} = \frac{\mu_0(\hat{\theta})\sum_{m\in M}\sigma(m|\hat{\theta})\phi(s'',t''|m)}{\mu_0(\bar{\theta})\sum_{m\in M}\sigma(m|\bar{\theta})\phi(s'',t''|m)}.$$

Thus,

$$\frac{\sum_{m\in M}\sigma(m|\hat{\theta})\phi(s,t|m)}{\sum_{m\in M}\sigma(m|\hat{\theta})\phi(s'',t''|m)} = \frac{\sum_{m\in M}\sigma(m|\bar{\theta})\phi(s,t|m)}{\sum_{m\in M}\sigma(m|\bar{\theta})\phi(s'',t''|m)}.$$

Since s'', t'' can be any one satisfying that $\mu(s'',t'') = \mu$, by the theorem on equal ratios, Equation 16 holds. $\square$

Sender's optmality and payoff equivalence: To show Sender's optimality and the payoff equivalence, it is sufficient to show that the expected payoffs of type θ Sender under both assessments are the same. The expected payoff of type θ Sender under the assessment (σ, α, μ) is

$$\sum_{m\in M,s\in S,t\geq 0,a\in A}\sigma(m|\theta)\phi(s,t|m)\alpha(a|s,t)(v(a)-t).$$

The expected payoff of type θ Sender under the assessment $(\sigma^*, \alpha^*, \mu^*)$ is

$$\sum_{\mu\in\mathbf{supp}\{\pi(\theta)\},a\in A}\pi(\mu|\theta)\alpha^*(a|\mu)(v(a)-x(\mu)).$$

These two expected payoffs are the same, since by Lemma 1, for any $a \in A, s \in S, t \geq 0$

the coefficient of $\alpha(a|s,t)$ where $\mu(s,t)=\mu\in\mathbf{supp}\{\pi(\theta)\}$ in the expression

$$\sum_{\mu\in\mathbf{supp}\{\pi(\theta)\},a\in A}\pi(\mu|\theta)\alpha^*(a|\mu)(v(a)-x(\mu))$$

is

$$\begin{aligned}
&\sum_{s',t',m,\mu(s',t')=\mu}\sigma(m|\theta)\phi(s',t'|m)\frac{\sum_{\theta'\in\Theta,m\in M}\mu_0(\theta')\sigma(m|\theta')\phi(s,t|m)}{\sum_{s'\in S,t'\geq 0,\mu(s',t')=\mu}\sum_{\theta'\in\Theta,m\in M}\mu_0(\theta')\sigma(m|\theta')\phi(s',t'|m)}v(a)\\
&=\sum_{s',t',m,\mu(s',t')=\mu}\sigma(m|\theta)\phi(s',t'|m)\frac{\sum_m\sigma(m|\theta)\phi(s,t|m)}{\sum_{s',t',m,\mu(s',t')=\mu}\sigma(m|\theta)\phi(s',t'|m)}v(a)\\
&=\sum_m\sigma(m|\theta)\phi(s,t|m)v(a),
\end{aligned}$$

and expected money burning of type θ Sender of $(\sigma^*,\alpha^*,\mu^*)$ is

$$\begin{aligned}
\sum_{\mu\in\mathbf{supp}\{\pi(\theta)\},a\in A}\pi(\mu|\theta)\alpha^*(a|\mu)x(\mu)&=\sum_{\mu\in\mathbf{supp}\{\pi(\theta)\}}\pi(\mu|\theta)x(\mu)\\
&=\sum_{\mu\in\mathbf{supp}\{\pi(\theta)\}}\pi(\mu|\theta)\frac{\sum_{s,t,\mu(s,t)=\mu,m\in M}\sigma(m|\theta)\phi(s,t|m)t}{\sum_{\mu(s,t)=\mu}\sum_{m\in M}\sigma(m|\theta)\phi(s,t|m)}\\
&=\sum_{m\in M,s\in S,t\geq 0}\sigma(m|\theta)\phi(s,t|m)t,
\end{aligned}$$

where the second equation is by Lemma 1.

Receiver's optmality: Since $\alpha(s,t)$ is the best response under the belief $\mu(s,t)$ and $\alpha^*(\mu)$ is a convex combination of some $\alpha(s',t')$ where $\mu(s',t')=\mu$, by the convexity of best response set, $\alpha^*(\mu)$ must satisfies the Receiver's optimality condition.

Bayes updating: Given $\mu\in\Delta(\Theta)$, for any s,t such that $\mu(s,t)=\mu=\mu^*(\mu)$, by Bayes updating, we have that for any $\theta\in\Theta$

$$\mu(\theta|s,t)\sum_{\theta'\in\Theta,m\in M}\mu_0(\theta')\sigma(m|\theta')\phi(s,t|m)=\mu_0(\theta)\sum_{m\in M}\sigma(m|\theta)\phi(s,t|m).$$

Hence,

$$\mu(\theta)\sum_{s,t,\mu(s,t)=\mu}\sum_{\theta'\in\Theta,m\in M}\mu_0(\theta')\sigma(m|\theta')\phi(s,t|m)=\sum_{s,t,\mu(s,t)=\mu}\mu_0(\theta)\sum_{m\in M}\sigma(m|\theta)\phi(s,t|m).$$

That is

$$\mu^*(\theta|\mu)\sum_{\theta'}\mu_0(\theta')\pi(\mu|\theta')=\mu_0(\theta)\pi(\mu|\theta).$$

$\square$

Proof of Proposition 2. According to belief-based approach, $p\in\Delta(\Delta(\Theta))$ and ex-post payoff V is induced by a canonical MDMB, if and only if they satisfy that $V(\mu)\in\mathbb{V}(\mu)$, $p\in BP(\mu_0)$ and for any $\theta,\theta'\in\Theta$,

$$\int_\mu \frac{\mu(\theta)}{\mu_0(\theta)}(V(\mu)-x(\mu))dp(\mu)\geq\int_\mu \frac{\mu(\theta')}{\mu_0(\theta')}(V(\mu)-x(\mu))dp(\mu).$$

By swap θ,θ' in above inequality, we can get

$$\int_\mu \frac{\mu(\theta)}{\mu_0(\theta)}(V(\mu)-x(\mu))dp(\mu)=\int_\mu \frac{\mu(\theta')}{\mu_0(\theta')}(V(\mu)-x(\mu))dp(\mu).$$

$\square$

Proof of Corollary 1. According to Proposition 2, we only need to show that, to calculate $\mathcal{V}^*(\mu_0)$, it is without loss of generality to assume $V(\mu)=\max\mathbb{V}(\mu)$.

Suppose there is $p\in BP(\mu_0),x,V:\Delta(\Theta)\to\mathbb{R}$ such that there exists $\mu\in\textbf{supp}\{p\}$ satisfying that $V(\mu)\neq\max\mathbb{V}(\mu)$, then we construct V',x' such that at posterior μ, $V'(\mu)=\max\mathbb{V}(\mu)$ and $x'(\mu)=x(\mu)+V'(\mu)-V(\mu)$. Now p,x',V' also satisfies the constraints of Equation 8 without reducing the Sender's payoff. $\square$

B.2 Omitted Proofs in Section 4

Proof of Proposition 3. Before delving into the main body of the proof, we first prove a lemma of upper bound of $\mathcal{V}^*(\mu_0)$ and show our construction method.

Lemma 2. $\mathcal{V}^*(\mu_0)\leq\max_\pi\min_{\theta\in\Theta}V_\pi(\theta)$.

Proof of Lemma 2. For any canonical MDMB (π,t), by incentive-compatible constraints, the Sender's payoff under this mechanism is

$$V_\pi(\theta)-\sum_{\mu\in\textbf{supp}\{\pi(\theta)\}}\pi(\mu|\theta)x(\mu),$$

and for any θ, θ',

$$V_\pi(\theta) - \sum_{\mu \in \mathbf{supp}\{\pi(\theta)\}} \pi(\mu|\theta)x(\mu) = V_\pi(\theta') - \sum_{\mu \in \mathbf{supp}\{\pi(\theta)\}} \pi(\mu|\theta')x(\mu).$$

Then by $x(m) \geq 0$, we know that the Sender's payoff is no larger than $\min_{\theta\in\Theta} V_\pi(\theta)$. Hence, $\mathcal{V}^*(\mu_0) \leq \max_\pi \min_{\theta\in\Theta} V_\pi(\theta)$. $\square$

In order to achieve the upper bound, we use the method implied by the following proposition to obtain the minimum interim signaling payoff for any given signaling scheme π.

Proposition 11. *Given any signaling scheme $\pi : \Theta \to \Delta(\Delta(\Theta))$, we construct an MDMB $(\bar{\pi}, \bar{x})$,*

$$\bar{\pi}(\mu|\theta) = \begin{cases} (1-\delta)\pi(\mu|\theta) & \mu(\theta|\theta) \neq 1 \\ \delta + (1-\delta)\pi(\mu|\theta) & \mu(\theta|\theta) = 1 \end{cases}, x(\cdot) = \begin{cases} 0 & \mu(\theta|\theta) \neq 1 \\ \frac{1}{\bar{\pi}(\mu|\theta)}(V_{\bar{\pi}}(\theta) - \min_{\theta\in\Theta} V_{\bar{\pi}}(\theta)) & \mu(\theta|\theta) = 1 \end{cases}.$$

The mechanism $(\bar{\pi}, \bar{x})$ is incentive-compatible for $\delta \in (0,1)$ and the Sender's payoff under this mechanism converges to $\min_{\theta\in\Theta} V_\pi(\theta)$ as $\delta \to 0^+$.

Proof of Proposition 11. Given that $V_{\bar{\pi}}(\theta) - \min_{\theta\in\Theta} V_{\bar{\pi}}(\theta) \geq 0$, it follows that $x(\mu) \geq 0$ for all $\mu \in \Delta(\Theta)$. Therefore, for any $\theta \in \Theta$, the Sender's payoff of type θ under the mechanism $(\bar{\pi}, x)$ is equal to

$$V_{\bar{\pi}}(\theta) - \sum_{\mu \in \mathbf{supp}\{\pi(\theta)\}} \bar{\pi}(\mu|\theta)x(\mu) = \min_{\theta\in\Theta} V_{\bar{\pi}}(\theta).$$

This implies that the mechanism is incentive-compatible. Moreover, the Sender's expected payoff is $\min_{\theta\in\Theta} V_{\bar{\pi}}(\theta)$.

Next, we show that the corresponding canonical assessments also satisfies Bayes updating condition under this MDMB mechanism. Let $\mu_\theta \in \Delta(\Theta)$ be such that $\mu_\theta(\theta) = 1$ and $\mu_\theta(\theta') = 0$ for any $\theta' \neq \theta$. Let $\mu^*_\pi(\mu)$ and $\mu^*_{\bar{\pi}}(\mu)$ denote the posterior beliefs under the signaling schemes π and $\bar{\pi}$ for any $\mu \in \Delta(\Theta)$, respectively. For any $\mu \in \Delta(\Theta)$, we have that

$$\mu^*_{\bar{\pi}}(\theta|\mu) = \frac{\bar{\pi}(\mu|\theta)\mu_0(\theta)}{\sum_{\theta'\in\Theta} \bar{\pi}(\mu|\theta')\mu_0(\theta')} = \frac{\pi(\mu|\theta)\mu_0(\theta)}{\sum_{\theta'\in\Theta} \pi(\mu|\theta')\mu_0(\theta')} = \mu^*_\pi(\theta|\mu).$$

Thus, $\mu^*_\pi(\mu) = \mu$ if and only if $\mu^*_{\bar{\pi}}(\mu) = \mu$. Hence, $V_{\bar{\pi}}(\theta) = \sum_{\mu \in \mathbf{supp}\{\pi(\theta)\}} \bar{\pi}(\mu|\theta)V(\mu) = (1-\delta)V_\pi(\theta) + \delta V(\mu_\theta)$. It follows that

$$\lim_{\delta \to 0^+} \min_{\theta \in \Theta} \{V_{\bar{\pi}}(\theta)\} = \lim_{\delta \to 0^+} \min_{\theta \in \Theta} \{(1-\delta)V_\pi(\theta) + \delta V(\mu_\theta)\} = \min_{\theta \in \Theta} V_\pi(\theta).$$

□

We now return to the problem of attaining the upper bound of the value of the MDMB derived by Lemma 2. As per the construction in Proposition 11, if the Sender adopts the signaling scheme that maximizes the worst-case interim signaling payoff, the Sender's payoff will achieve $\max_\pi \min_{\theta \in \Theta} V_\pi(\theta)$. Hence, we obtain $\mathcal{V}^*(\mu_0) = \max_\pi \min_{\theta \in \Theta} V_\pi(\theta)$. So, we obtain the upper bound and get that $\mathcal{V}^*(\mu_0) = \max_\pi \min_{\theta \in \Theta} V_\pi(\theta)$.

For $\lambda \in \Delta(\Theta), p \in BP(\mu_0)$, let

$$\mathcal{L}(\lambda, p) \triangleq \int_\mu \hat{V}_\lambda(\mu) dp(\mu).$$

We have $\Delta(\Theta), BP(\mu_0)$ are both convex and compact. Further by $\hat{V}(\mu)$ is upper semi-continuous [14], then $\mathcal{L}(\lambda, p)$ is continuous and linear in λ, and also it is upper semi-continuous and linear in p. Thus according to Sion's minimax theorem, we have that

$$\max_{p \in BP(\mu_0)} \min_{\lambda \in \Delta(\Theta)} \mathcal{L}(\lambda, p) = \min_{\lambda \in \Delta(\Theta)} \max_{p \in BP(\mu_0)} \mathcal{L}(\lambda, p).$$

Next by Kamenica and Gentzkow [2011], we know that for any given λ,

$$\max_{p \in BP(\mu_0)} \mathcal{L}(\lambda, p) = cav(\hat{V}_\lambda)(\mu_0).$$

Since $\mathcal{L}(\lambda, p)$ is linear in λ, then

$$\max_{p \in BP(\mu_0)} \min_{\lambda \in \Delta(\Theta)} \mathcal{L}(\lambda, p) = \max_\pi \min_{\theta \in \Theta} V_\pi(\theta) = \mathcal{V}^*(\mu_0).$$

Therefore, $\mathcal{V}^*(\mu_0) = \min_{\lambda \in \Delta(\Theta)} cav(\hat{V}_\lambda)(\mu_0)$.

□

Proof of Proposition 4. Since $\Delta(\Theta)$ and $BP(\mu_0)$ are compact sets, by extensions of Sion's

[14]See Lipnowski and Ravid [2020] footnote 13

minmax theorem Arandjelović [1992], there exists a saddle point (λ_0, p_0) such that

$$\max_{p\in BP(\mu_0)} \min_{\lambda\in\Delta(\Theta)} \mathcal{L}(\lambda, p) = \min_{\lambda\in\Delta(\Theta)} \max_{p\in BP(\mu_0)} \mathcal{L}(\lambda, p) = \mathcal{L}(\lambda_0, p_0).$$

Next, we show that for any saddle point (λ_1, p_1) of $\mathcal{L}$, we have $\mathcal{L}(\lambda_0, p_0) = \mathcal{L}(\lambda_1, p_1)$. We prove by contradiction and assume, without loss of generality, that $\mathcal{L}(\lambda_0, p_0) < \mathcal{L}(\lambda_1, p_1)$. Then, by the property of saddle point, we have

$$\mathcal{L}(\lambda_0, p_1) \leq \mathcal{L}(\lambda_0, p_0) < \mathcal{L}(\lambda_1, p_1).$$

The inequality $\mathcal{L}(\lambda_0, p_1) < \mathcal{L}(\lambda_1, p_1)$ contradicts the fact that (λ_1, p_1) is a saddle point.

Suppose that λ^* satisfies the following conditions: there exists $p^* \in BP(\mu_0)$ such that $\mathcal{L}(\lambda^*, p^*) = cav(\hat{V}_{\lambda^*})(\mu_0)$, and for any $\theta \in \mathbf{supp}(\lambda^*)$, $\mathcal{L}(\lambda^*, p^*) = \mathcal{L}(\theta, p^*)$ and for any $\theta \notin \mathbf{supp}(\lambda^*)$, $\mathcal{L}(\lambda^*, p^*) \leq \mathcal{L}(\theta, p^*)$. Then, by definition, we have that (λ^*, p^*) is a saddle point and $\mathcal{L}(\lambda^*, p^*) = \mathcal{L}(\lambda_0, p_0)$. Hence, λ^* is the worst subjective prior.

If λ^* is the worst subjective prior, then choosing p^* is optimal and we can obtain that

$$\mathcal{L}(\lambda_0, p_0) = \min_{\lambda} \max_{p} \mathcal{L}(\lambda, p) \geq \mathcal{L}(\lambda^*, p^*) \geq \max_{p} \min_{\lambda} \mathcal{L}(\lambda, p) = \mathcal{L}(\lambda_0, p_0).$$

Therefore, $\mathcal{L}(\lambda^*, p^*) = \mathcal{L}(\lambda_0, p_0)$ is the mini-max or max-min value. It follows that for any $\lambda \in \Delta(\Theta)$, $\mathcal{L}(\lambda, p^*) \geq \mathcal{L}(\lambda^*, p^*)$ and for any $p \in BP(\mu_0)$, $\mathcal{L}(\lambda^*, p) \leq \mathcal{L}(\lambda^*, p^*)$. Thus, (λ^*, p^*) is a saddle point and it must satisfy that $\mathcal{L}(\lambda^*, p^*) = cav(\hat{V}_{\lambda^*})(\mu_0)$, and for any $\theta \in \mathbf{supp}(\lambda^*)$, $\mathcal{L}(\lambda^*, p^*) = \mathcal{L}(\theta, p^*)$ and for any $\theta \notin \mathbf{supp}(\lambda^*)$, $\mathcal{L}(\lambda^*, p^*) \leq \mathcal{L}(\theta, p^*)$. □

Proof of Proposition 5. The proof is analogous to the proof of Proposition 4. □

Proof of Proposition 6. By Proposition 4, it suffices to show that there exists $\theta_i \in \{\theta_1, \theta_2\}$ such that for any $p \in BP(\mu_0)$ with $\mathcal{L}(\theta_i, p) = cav(\hat{V}_{\theta_i})(\mu_0)$, we have $\mathcal{L}(\theta_i, p) \leq \mathcal{L}(\theta_{3-i}, p)$.

Let $U = \{\mu | V(\mu) = \max_{x\in[0,1]} V(x)\}$ denote the range of posteriors that yield the maximum value for Sender, for $\mu \in [0, 1]$. Since U is convex and $V(\cdot)$ is upper semi-continuous, U can be expressed as the union of closed intervals. We assume that $l = \min U$ and $r = \max U$. If $l \leq \mu_0 \leq r$, then it is clear that $V^*(\mu_0) = \max_{x\in[0,1]} V(x)$ and for any θ_i, $cav(\hat{V}_{\theta_i})(\mu_0) = V^*(\mu_0)$, which is a constant. Hence, our statement holds trivially. In the following proof, we consider $\mu_0 > r$ or $\mu_0 < l$.

Without loss of generality, we assume that $\mu_0 > r \geq 0$ by symmetry. We focus on θ_1. We prove by contradiction that if there exists $p \in BP(\mu_0)$ with $\mathcal{L}(\theta_1, p) = cav(\hat{V}_{\theta_1})(\mu_0)$ and $\mathcal{L}(\theta_1, p) > \mathcal{L}(\theta_2, p)$, then we reach a contradiction. Since $p \in BP(\mu_0)$ performs the concavification of the function $\hat{V}_{\theta_1}$ at point μ_0, by Proposition 9 of the working paper version of Kamenica and Gentzkow [2011], we have that the points $(\mu, \hat{V}_{\theta_1}(\mu))$ for $\mu \in \textbf{supp}\{p\}$ are collinear. This means that there exist parameters k, b such that for any $\mu \in \textbf{supp}\{p\}$,

$$\frac{\mu}{\mu_0} V(\mu) = k\mu + b.$$

Since $\mathcal{L}(\theta_1, p) > \mathcal{L}(\theta_2, p)$, we have

$$\int_\mu (\mu - \mu_0) V(\mu) dp(\mu) > 0.$$

Substituting $V(\mu)$ with $\mu_0(k + \frac{b}{\mu})$ and using $\int_\mu (\mu - \mu_0) dp(\mu) = 0$, we obtain

$$b(1 - \int_\mu \frac{\mu_0}{\mu} dp(\mu)) = \int_\mu (\mu - \mu_0) \frac{b}{\mu} dp(\mu) > 0.$$

By Cauchy's inequality,

$$\int_\mu \frac{\mu_0}{\mu} dp(\mu) = \int_\mu \frac{\mu}{\mu_0} dp(\mu) \int_\mu \frac{\mu_0}{\mu} dp(\mu) \geq (\int_\mu dp(\mu))^2 = 1.$$

Therefore, we must have $b < 0$. However, $k\mu + b$ is the concavification line of $\hat{V}_{\theta_1}(\cdot)$ at μ_0. Thus, it must satisfy that for any $\mu \in [0, 1]$, $\hat{V}_{\theta_1}(\mu) \leq k\mu + b$. Choosing $\mu = 0$, we get $b \geq 0$. This is a contradiction. □

B.3 Omitted Proofs in Section 5

Proof of Proposition 7. To facilitate the whole proof, we first establish two lemmas that reveal some useful properties of $\mathcal{V}^*(\mu)$ and one important proposition.

Lemma 3. *For any signaling scheme $p \in BP(\mu_0)$, we have that*

$$\mathcal{V}^*(\mu_0) \geq \min_{\theta \in \textbf{supp}\{\mu_0\}} \int_\mu \frac{\mu(\theta)}{\mu_0(\theta)} \mathcal{V}^*(\mu) dp(\mu).$$

Proof of Lemma 3. Let p_μ denote the optimal signaling scheme that maximizes the minimum interim payoff under any prior $\mu \in \mathbf{supp}\{p\}$. We consider a signaling scheme for μ_0 that consists of two stages: first, it splits μ_0 according to p, and then, for each $\mu \in \mathbf{supp}\{p\}$, it further splits μ according to p_μ. By applying this scheme, we obtain the following inequality:

$$\mathcal{V}^*(\mu_0) \geq \min_{\theta \in \mathbf{supp}\{\mu_0\}} \{\int_\mu \int_{\mu'} \frac{\mu'(\theta)}{\mu_0(\theta)} V(\mu') dp(\mu) dp_\mu(\mu')\} = \min_{\theta \in \mathbf{supp}\{\mu_0\}} \int_\mu \frac{\mu(\theta)}{\mu_0(\theta)} \mathcal{V}^*(\mu) dp(\mu).$$

□

Lemma 4. *$\mathcal{V}^*(\mu)$ is continuous at any full-support point μ, i.e. $\mu(\theta) > 0$ for all $\theta \in \Theta$.*

Proof of Lemma 4. We begin by recalling that, by Proposition 3, $\mathcal{V}^*(\mu)$ is the infimum of a family of continuous functions, namely the concave closures. Hence, $\mathcal{V}^*(\mu)$ is upper semi-continuous. Without loss of generality, we assume that $v(a) \geq 0$ for all $a \in A$. Suppose that $\mathcal{V}^*(\mu)$ is discontinuous at some point μ with full support. Then there exist two beliefs $\mu_1, \mu_2 \in \Delta(\Theta)$ and a positive constant D such that, for some sufficiently small constant $B(D, \mu_1, \mu_2)$ (denoted by B), we have $\mathcal{V}^*(\mu) > \mathcal{V}^*(\mu + \varepsilon(\mu_1 - \mu_2)) + D$ for all $\varepsilon \in (0, B)$. Let $\mu_\varepsilon = \mu + \varepsilon(\mu_1 - \mu_2)$ for any $\varepsilon > 0$ such that $\varepsilon < B$ and $\varepsilon^2 < B$. We can rewrite μ_ε as $\mu_\varepsilon = (1 - \sqrt{\varepsilon})\mu + \sqrt{\varepsilon}(\mu + \sqrt{\varepsilon}(\mu_1 - \mu_2))$. Applying Lemma 3, we get

$$\begin{aligned}\mathcal{V}^*(\mu_\varepsilon) &\geq \min_\theta\{(1 - \sqrt{\varepsilon}) \frac{\mu(\theta)}{\mu_\varepsilon(\theta)} \mathcal{V}^*(\mu) + \sqrt{\varepsilon} \frac{\mu(\theta) + \sqrt{\varepsilon}(\mu_1(\theta) - \mu_2(\theta))}{\mu_\varepsilon(\theta)} \mathcal{V}^*(\mu + \sqrt{\varepsilon}(\mu_1 - \mu_2))\} \\ &\geq \min_\theta\{(1 - \sqrt{\varepsilon}) \frac{\mu(\theta)}{\mu_\varepsilon(\theta)} \mathcal{V}^*(\mu)\}.\end{aligned}$$

Because $\mu(\theta) > 0$ for any θ, $\lim_{\varepsilon \to 0^+}(1 - \sqrt{\varepsilon}) \frac{\mu(\theta)}{\mu_\varepsilon(\theta)} = 1$. Therefore it follows that $\lim_{\varepsilon \to 0^+} \mathcal{V}^*(\mu_\varepsilon) \geq \mathcal{V}^*(\mu)$, which contradicts the assumption that $\mathcal{V}^*(\mu_\varepsilon) < \mathcal{V}^*(\mu) - D$. □

Proposition 12. *Assuming that $\mathcal{V}^*(\mu)$ is continuous, we obtain $\mathcal{V}^*_{CT}(\mu_0) = \mathcal{V}^*_{MD}(\mu_0) = \mathcal{V}^*(\mu_0)$ when $\mathcal{V}^*(\mu_0) = \mathcal{V}^*_{MD}(\mu_0)$.*

Proof of Proposition 12. We adopt a proof by contradiction. Let μ_0 be a belief with the *smallest* support such that $\mathcal{V}^*_{CT}(\mu_0) < \mathcal{V}^*(\mu_0) = \mathcal{V}^*_{MD}(\mu_0)$. Then, applying Proposition 3 and Proposition 10, we infer that there is a posterior distribution p^* that yields the same interim payoff for every θ under p^* and this payoff is $\mathcal{V}^*(\mu_0)$. Therefore, p^* is a optimal signaling scheme defined in Section 4.1. By Proposition 5, we obtain that there is a worst subjective prior $\lambda^* \in \Delta(\Theta)$ such that for any θ,

$$\int_\mu \frac{\mu(\theta)}{\mu_0(\theta)} V(\mu) dp^*(\mu) = \int_\mu \sum_\theta \frac{\lambda^*(\theta)\mu(\theta)}{\mu_0(\theta)} V(\mu) dp^*(\mu),$$

and p^* is the concavification of $\sum_\theta \frac{\lambda^*(\theta)\mu(\theta)}{\mu_0(\theta)} V(\mu)$ at point μ_0. Then, by Proposition 9 of the working paper version of Kamenica and Gentzkow [2011], we have that there exist parameters A_θ for $\theta \in \Theta$ and for any $\mu \in \mathbf{supp}\{p^*\}$,

$$\sum_\theta \frac{\lambda^*(\theta)\mu(\theta)}{\mu_0(\theta)} V(\mu) = \sum_\theta A_\theta \mu(\theta).$$

Let $\Theta' = \{\theta \big| \lambda^*(\theta) = 0\}$, $U_2 = \{\mu \in \mathbf{supp}\{p^*\} \big| \mathbf{supp}\{\mu\} \subseteq \Theta'\}$, and $U_1 = \mathbf{supp}\{p^*\}/U_2$.

Hence, we obtain that for any $\theta \in \Theta$,

$$A_\theta \mu_0(\theta) \int_\mu \frac{\mu(\theta)}{\mu_0(\theta)} V(\mu) dp^*(\mu) = A_\theta \mu_0(\theta) \sum_{\theta'} A_{\theta'} \mu_0(\theta').$$

Summing them up and since $0 = \sum_{\theta\in\Theta} \lambda^*(\theta)\mu(\theta) = \sum_{\theta\in\Theta} A_\theta \mu(\theta)$ for any $\mu \in U_2$, we have that

$$\int_{\mu\in U_1} \sum_\theta A_\theta \mu(\theta) \frac{\sum_\theta A_\theta \mu(\theta)}{\sum_\theta \frac{\lambda^*(\theta)\mu(\theta)}{\mu_0(\theta)}} dp^*(\mu) = (\sum_\theta A_\theta \mu_0(\theta))^2.$$

Since $\int_{\mu\in U_1} \sum_\theta \frac{\lambda^*(\theta)\mu(\theta)}{\mu_0(\theta)} dp^*(\mu) = 1$, by Cauchy's inequality and the above equation, we also have that for any $\mu, \mu' \in U_1$,

$$\frac{\sum_\theta A_\theta \mu(\theta)}{\sum_\theta \frac{\lambda^*(\theta)\mu(\theta)}{\mu_0(\theta)}} = \frac{\sum_\theta A_\theta \mu'(\theta)}{\sum_\theta \frac{\lambda^*(\theta)\mu'(\theta)}{\mu_0(\theta)}}.$$

This implies that $V(\mu) = V(\mu')$ for $\mu, \mu' \in U_1$. To simplify the notation, we set this value to be R, which coincides with $\mathcal{V}^*(\mu_0)$ and $\mathcal{V}^*_{MD}(\mu_0)$.

If $U_2 = \emptyset$, then it is also a cheap talk equilibrium and we obtain that $\mathcal{V}^*(\mu_0) = \mathcal{V}^*_{CT}(\mu_0)$. Otherwise, let $p_1 = \int_{\mu\in U_1} dp^*(\mu), p_2 = \int_{\mu\in U_2} dp^*(\mu)$ and $\mu_1 = \int_{\mu\in U_1} \mu dp^*(\mu)/p_1, \mu_2 = \int_{\mu\in U_2} \mu dp^*(\mu)/p_2$. Since $\mathbf{supp}\{\mu_2\} < \mathbf{supp}\{\mu_0\}$, we must have $\mathcal{V}^*(\mu_2) = \mathcal{V}^*_{MD}(\mu_2)$ if and only if $\mathcal{V}^*(\mu_2) = \mathcal{V}^*_{CT}(\mu_2)$. Then for any $\theta \in \mathbf{supp}\{\mu_2\}$, we have that

$$R = \int_\mu \frac{\mu(\theta)}{\mu_0(\theta)} V(\mu) dp^*(\mu) = p_1 \frac{\mu_1(\theta)}{\mu_0(\theta)} R + p_2 \frac{\mu_2(\theta)}{\mu_0(\theta)} \int_{\mu\in U_2} \frac{\mu(\theta)}{\mu_2(\theta)} V(\mu) d\frac{p^*(\mu)}{p_2}.$$

Hence, for any $\theta \in \textbf{supp}\{\mu_2\}$,

$$\int_{\mu \in U_2} \frac{\mu(\theta)}{\mu_2(\theta)} V(\mu) d\frac{p^*(\mu)}{p_2} = R.$$

This implies that $\mathcal{V}^*_{MD}(\mu_2) \geq R$. Next we will divide our final proof into two cases.

Case 1: If $\mathcal{V}^*(\mu_2) = R$, we have $\mathcal{V}^*(\mu_2) = \mathcal{V}^*_{MD}(\mu_2)$, which implies that $\mathcal{V}_{CT}(\mu_2) = R$ by the smallest support property of μ_0. Since $V(\mu) = R$ for all $\mu \in U_1$, we can deduce that there exists a cheap talk equilibrium for μ_0 where Sender achieves payoff R. This means that we have shown that $\mathcal{V}_{CT}(\mu_0) = R$, which contradicts our assumption.

Case 2: Suppose that $\mathcal{V}^*(\mu_2) > R$. Given the continuity of $\mathcal{V}^*(\mu)$, we can find a positive constant $\varepsilon > 0$, such that $\mathcal{V}^*(\frac{\varepsilon\mu_1 + p_2\mu_2}{\varepsilon + p_2}) > R$ and $\mu_0 = (p_1 - \varepsilon)\mu_1 + (\varepsilon + p_2)\frac{\varepsilon\mu_1 + p_2\mu_2}{\varepsilon + p_2}$. Hence by Lemma 3, we have that

$$\mathcal{V}^*(\mu_0) \geq \min_{\theta \in \textbf{supp}\{\mu_0\}} \{(p_1 - \varepsilon)\frac{\mu_1(\theta)}{\mu_0(\theta)} \mathcal{V}^*(\mu_1) + \frac{\varepsilon\mu_1(\theta) + p_2\mu_2(\theta)}{\mu_0(\theta)} \mathcal{V}^*(\frac{\varepsilon\mu_1 + p_2\mu_2}{\varepsilon + p_2})\}.$$

Since $\mathcal{V}^*(\mu_1) \geq R$, $\mathcal{V}^*(\frac{\varepsilon\mu_1 + p_2\mu_2}{\varepsilon + p_2}) > R$ and for $\theta \in \textbf{supp}\{\mu_0\}$, $\varepsilon\mu_1(\theta) + p_2\mu_2(\theta) > 0$, we will get that

$$\mathcal{V}^*(\mu_0) > R,$$

which is also a contradiction. □

Back to the proof of Proposition 7, without loss of generality, we assume that $v(a) \geq 0$ for all $a \in A$. According to Proposition 12 and Lemma 4, it is suffice to show that under generic settings $\mathcal{V}^*$ is continuous at the boundary of $\Delta(\Theta)$. To prove this, we firstly prove the following lemma.

Lemma 5. *For any belief μ at the boundary and any direction $\mu_1 \in \Delta(\Theta)$, we have*

$$\lim_{\varepsilon \to 0^+} \mathcal{V}^*((1-\varepsilon)\mu + \varepsilon\mu_1) \geq V(\mu).$$

proof of Lemma 5. If $\textbf{supp}\{\mu_1\} \subseteq \textbf{supp}\{\mu\}$, then by Lemma 4, we can get

$$\lim_{\varepsilon \to 0^+} \mathcal{V}^*((1-\varepsilon)\mu + \varepsilon\mu_1) = \mathcal{V}^*(\mu) \geq V(\mu).$$

Consider the case where $\textbf{supp}\{\mu_1\} \not\subseteq \textbf{supp}\{\mu\}$. By generic setting Definition 3, we can find a belief μ_2 such that $\mathbf{supp}\{\mu_2\} = \mathbf{supp}\{\mu\}$, $V(\mu_2) = V(\mu)$, and the singleton $RO(\mu_2) \subseteq RO(\mu)$. It follows that for any $\gamma \in (0,1)$, we have $RO((1-\gamma)\mu+\gamma\mu_2) = RO(\mu_2)$ and $V((1-\gamma)\mu+\gamma\mu_2) = V(\mu)$. Let $\mu_\gamma = (1-\gamma)\mu+\gamma\mu_2$. Since $RO(\mu_\gamma)$ is a singleton, implying that the action in this singleton dominates all other actions, we select ε_γ such that $0 < \varepsilon_\gamma < 1$ and $RO(\mu_2) \subseteq RO((1-\varepsilon_\gamma)\mu_\gamma + \varepsilon_\gamma\mu_1)$. We then proceed to examine two cases.

Case 1: $\lim_{\gamma\to 0^+} \varepsilon_\gamma \neq 0$, which means that $\lim_{\varepsilon\to 0^+} V(\varepsilon\mu_1 + (1-\varepsilon)\mu) = V(\mu)$. Thus,

$$\lim_{\varepsilon\to 0^+} \mathcal{V}^*((1-\varepsilon)\mu + \varepsilon\mu_1) \geq \lim_{\varepsilon\to 0^+} V((1-\varepsilon)\mu + \varepsilon\mu_1) = V(\mu).$$

Case 2: $\lim_{\gamma\to 0^+} \varepsilon_\gamma = 0$. Then consider following splitting scheme of $(1-\varepsilon_\gamma)\mu + \varepsilon_\gamma\mu_1$ where

$$(1-\varepsilon_\gamma)\mu + \varepsilon_\gamma\mu_1 = \min_\theta \frac{(1-\varepsilon_\gamma)\mu(\theta) + \varepsilon_\gamma\mu_1(\theta)}{(1-\varepsilon_\gamma)\mu_\gamma(\theta) + \varepsilon_\gamma\mu_1(\theta)}((1-\varepsilon_\gamma)\mu_\gamma + \varepsilon_\gamma\mu_1) + P_2\mu'$$

where $P_2 = 1 - \min_\theta \frac{(1-\varepsilon_\gamma)\mu(\theta)+\varepsilon_\gamma\mu_1(\theta)}{(1-\varepsilon_\gamma)\mu_\gamma(\theta)+\varepsilon_\gamma\mu_1(\theta)}$ and $\mu' \in \Delta(\Theta)$. By Lemma 3 we have that

$$\begin{aligned}\lim_{\varepsilon\to 0^+} \mathcal{V}^*((1-\varepsilon)\mu + \varepsilon\mu_1) &= \lim_{\gamma\to 0^+} \mathcal{V}^*((1-\varepsilon_\gamma)\mu + \varepsilon_\gamma\mu_1) \\ &\geq \lim_{\gamma\to 0^+} \min_\theta \frac{(1-\varepsilon_\gamma)\mu(\theta) + \varepsilon_\gamma\mu_1(\theta)}{(1-\varepsilon_\gamma)\mu_\gamma(\theta) + \varepsilon_\gamma\mu_1(\theta)} \min_\theta \frac{(1-\varepsilon_\gamma)\mu_\gamma(\theta) + \varepsilon_\gamma\mu_1(\theta)}{(1-\varepsilon_\gamma)\mu(\theta) + \varepsilon_\gamma\mu_1(\theta)} V(\mu).\end{aligned}$$

For any $\theta \in \mathbf{supp}\{\mu\}$,

$$\lim_{\gamma\to 0^+} \frac{(1-\varepsilon_\gamma)\mu(\theta) + \varepsilon_\gamma\mu_1(\theta)}{(1-\varepsilon_\gamma)\mu_\gamma(\theta) + \varepsilon_\gamma\mu_1(\theta)} = \lim_{\gamma\to 0^+} \frac{\mu(\theta)}{\mu_\gamma(\theta)} = 1,$$

and for any $\theta \in \mathbf{supp}\{\mu_1\}/\mathbf{supp}\{\mu\}$,

$$\lim_{\gamma\to 0^+} \frac{(1-\varepsilon_\gamma)\mu_0(\theta) + \varepsilon_\gamma\mu_1(\theta)}{(1-\varepsilon_\gamma)\mu_\gamma(\theta) + \varepsilon_\gamma\mu_1(\theta)} = \lim_{\gamma\to 0^+} \frac{\varepsilon_\gamma\mu_1(\theta)}{\varepsilon_\gamma\mu_1(\theta)} = 1.$$

So we obtain that for any $\mu_1 \in \Delta(\Theta)$,

$$\lim_{\varepsilon\to 0^+} \mathcal{V}^*((1-\varepsilon)\mu + \varepsilon\mu_1) \geq V(\mu).$$

□

Building on Lemma 5, we demonstrate the continuity of $\mathcal{V}^*$ at any point μ. Let p^* be the optimal signaling scheme of μ_0 that attains the minimum interim payoff, that is, $\mathcal{V}^*(\mu_0) = \min_{\theta\in\mathbf{supp}\{\mu_0\}} \int_\mu \frac{\mu(\theta)}{\mu_0(\theta)} V(\mu)dp^*(\mu)$. Then by Lemma 3, for any $\mu_1 \in \Delta(\Theta)$, we have that

$$\begin{aligned}\lim_{\varepsilon\to 0^+} \mathcal{V}^*((1-\varepsilon)\mu_0 + \varepsilon\mu_1) &\geq \lim_{\varepsilon\to 0^+} \min_\theta \int_\mu \frac{(1-\varepsilon)\mu(\theta)+\varepsilon\mu_1(\theta)}{(1-\varepsilon)\mu_0(\theta)+\varepsilon\mu_1(\theta)} \mathcal{V}^*((1-\varepsilon)\mu+\varepsilon\mu_1)dp^*(\mu)\\ &\geq \mathcal{V}^*(\mu_0).\end{aligned}$$

The last inequality holds because

$$\lim_{\varepsilon\to 0^+} \min_{\theta\in\mathbf{supp}\{\mu_0\}} \int_\mu \frac{(1-\varepsilon)\mu(\theta)+\varepsilon\mu_1(\theta)}{(1-\varepsilon)\mu_0(\theta)+\varepsilon\mu_1(\theta)} \mathcal{V}^*((1-\varepsilon)\mu+\varepsilon\mu_1)dp^*(\mu) \geq \min_{\theta\in\mathbf{supp}\{\mu_0\}} \int_\mu \frac{\mu(\theta)}{\mu_0(\theta)} V(\mu)dp^*(\mu),$$

and for $\theta \in \mathbf{supp}\{\mu_1\}/\mathbf{supp}\{\mu_0\}$

$$\begin{aligned}\lim_{\varepsilon\to 0^+} \int_\mu \frac{(1-\varepsilon)\mu(\theta)+\varepsilon\mu_1(\theta)}{(1-\varepsilon)\mu_0(\theta)+\varepsilon\mu_1(\theta)} \mathcal{V}^*((1-\varepsilon)\mu+\varepsilon\mu_1)dp^*(\mu) &\geq \int_\mu V(\mu)dp^*(\mu)\\ &= \sum_\theta \mu_0(\theta) \int_\mu \frac{\mu(\theta)}{\mu_0(\theta)} V(\mu)dp^*(\mu)\\ &\geq \min_{\theta\in\mathbf{supp}\{\mu_0\}} \int_\mu \frac{\mu(\theta)}{\mu_0(\theta)} V(\mu)dp^*(\mu).\end{aligned}$$

Using the fact that $\mathcal{V}^*$ is upper-semi continuous, we establish the continuity of $\mathcal{V}^*$. □

Proof of Proposition 8. By Proposition 3, we only need to construct a posterior distribution $p \in BP(\mu_0)$ such that the minimum interim payoff under this distribution is greater than $\mathcal{V}^*_{CT}(\mu_0)$.

Since the setting is generic, then there must be a full-support $\hat{\mu}$ such that $V(\hat{\mu}) = \max_{\mu\in\Delta(\Theta)} V(\mu)$. We define $\mu(x) = \frac{\mu_0 - x\hat{\mu}}{1-x}$ and, since μ_0 has full support, there exists a small enough $\varepsilon > 0$ such that $\mu(\varepsilon) \in \Delta(\Theta)$ and $qcav(V)(\mu(\varepsilon)) = qcav(V)(\mu_0)$. Let τ be the distribution of posterior that is the quasi-concavification of V at point $\mu(\varepsilon)$. Then we construct the distribution of posterior τ^* that acts as τ with probability $1-\varepsilon$ and induces the posterior $\hat{\mu}$ with probability ε. Then the interim payoff of type θ under τ^* is

$$qcav(V)(\mu_0) + \varepsilon \frac{\hat{\mu}(\theta)}{\mu_0(\theta)}(V(\hat{\mu}) - qcav(V)(\mu_0))$$

Since $qcav(V)(\mu_0) \neq cav(V)(\mu_0)$, we have that $V(\hat{\mu}) = \max_{\mu\in\Delta(\Theta)} V(\mu) > qcav(V)(\mu_0)$.

Hence, we obtain that $\mathcal{V}^*(\mu_0) > \mathcal{V}^*_{CT}(\mu_0)$.

□

B.4 Omitted Proofs in Section 6

Proof of Proposition 9. Since it is easy to obtain that $\mathcal{V}^*_{CT}(\mu_0) \leq \mathcal{V}^*_{MD}(\mu_0) \leq \mathcal{V}^*(\mu_0) \leq \mathcal{V}^*_{BP}(\mu_0)$, it suffices to show that $\mathcal{V}^*_{CT}(\mu_0) = \mathcal{V}^*_{BP}(\mu_0)$. Since $\mathcal{V}^*(\mu_0) = \mathcal{V}^*_{BP}(\mu_0)$, by Proposition 3, let $\lambda^* \in \Delta(\Theta)$ be the worst subjective prior and $p^* \in BP(\mu_0)$ be the optimal distribution of posterior that maximizes the minimum interim payoff. Then we have $\mathcal{L}(\lambda^*, p^*) = \mathcal{V}^*_{BP}(\mu_0)$. Hence, for any θ, we must have

$$\int_\mu \frac{\mu(\theta)}{\mu_0(\theta)} \max \mathbb{V}(\mu) dp^*(\mu) = \mathcal{V}^*_{BP}(\mu_0),$$

Otherwise, $\mathcal{L}(\lambda^*, p^*) = \min_\theta \int_\mu \frac{\mu(\theta)}{\mu_0(\theta)} \max \mathbb{V}(\mu) dp^*(\mu) < \mathcal{V}^*_{BP}(\mu_0)$, since $\mathcal{L}(\lambda^*, p^*)$ is the minimum interim payoff and $\mathcal{V}^*_{BP}(\mu_0)$ is the expected interim payoff. Therefore,

$$\mathcal{V}^*_{BP}(\mu_0) = \sum_\theta \mu_0(\theta) \mathcal{V}^*_{BP}(\mu_0) = \int_\mu \max \mathbb{V}(\mu) dp^*(\mu).$$

Thus, p^* is also the concavification of $\max \mathbb{V}(\cdot)$ at point μ_0. By the same technique of proposition 9 of the working paper version of Kamenica and Gentzkow [2011], we deduce that $(\mu, \max \mathbb{V}(\mu))$ for $\mu \in \mathbf{supp}\{p^*\}$ are coplanar. This means that there exist parameters A_θ for $\theta \in \Theta$ such that for $\mu \in \mathbf{supp}\{p^*\}$,

$$\max \mathbb{V}(\mu) = \sum_\theta A_\theta \mu(\theta).$$

Then, for any θ, we have

$$\int_\mu \mu(\theta) \sum_{\theta'} A_{\theta'} \mu(\theta') dp^*(\mu) = \mu_0(\theta) \sum_\theta A_\theta \mu_0(\theta),$$

Multiplying by A_θ and summing over θ, we obtain

$$\int_\mu (\sum_\theta A_\theta \mu(\theta))^2 dp^*(\mu) = (\sum_\theta A_\theta \mu_0(\theta))^2.$$

By Cauchy inequality, we have

$$(\sum_\theta A_\theta \mu_0(\theta))^2 = \int_\mu dp^*(\mu) \int_\mu (\sum_\theta A_\theta \mu(\theta))^2 dp^*(\mu) \geq (\int_\mu A_\theta \mu(\theta) dp^*(\mu))^2 = (\sum_\theta A_\theta \mu_0(\theta))^2.$$

Therefore, we get $\sum_\theta A_\theta \mu(\theta) = \sum_\theta A_\theta \mu'(\theta)$ for all $\mu, \mu' \in \mathbf{supp}\{p^*\}$, which implies that $\max \mathbb{V}(\mu) = \max \mathbb{V}(\mu')$. This means that all the posterior beliefs induce the same value for Sender, so p^* is also a cheap talk equilibrium, i.e., Sender cannot find a more profitable message. Hence, $\mathcal{V}^*_{CT}(\mu_0) = \mathcal{V}^*_{BP}(\mu_0)$. □

Proof of Proposition 10. In this proof, $V(\mu)$ is no longer simply represents $\max \mathbb{V}(\mu)$. Back to the proof of Proposition 1, our construction of x in Equation 15 satisfies that if it is only possible that $t \leq C$, then $x(\mu) \leq C$ for all μ. Hence, the revelation principle can directly be generalized to this case by adding a new constraint that $x(\mu) \leq C$ for all $\mu \in \Delta(\Theta)$.

According to Proposition 2, Corollary 1 and previous discussions, we begin by characterizing $\mathcal{V}^*_C(\mu_0)$ as following optimization problem.

$$\begin{aligned}
\max \quad & k && (17)\\
\text{s.t} \quad & k = \int_\mu \frac{\mu(\theta)}{\mu_0(\theta)}(V(\mu) - x(\mu))dp(\mu) \quad \forall \theta \in \Theta && \text{(IC)}\\
& p \in BP(\mu_0) && \text{(Bayes plausible)}\\
& V(\mu) \in \mathbb{V}(\mu) \quad \forall \mu \in \Delta(\Theta) && \text{(Obedience)}\\
& 0 \leq x(\mu) \leq C \quad \forall \mu \in \Delta(\Theta) &&
\end{aligned}$$

We adopt a two-step optimization approach. First, we fix the signaling scheme $p \in BP(\mu_0)$ and obedience condition $V(\mu)$ and then we find the optimal burning scheme $x(\mu)$ where $0 \leq x(\mu) \leq C$. Thus, now it is a linear programming problem. By the fundamental theorem of linear programming, we can also obtain $\mathcal{V}^*_C(\mu_0)$ from the following max-min problem.

$$\begin{aligned} \max_{p,V} \quad \min_{\lambda} \quad & \int_\mu (\sum_{\theta\in\Theta}\lambda(\theta)\frac{\mu(\theta)}{\mu_0(\theta)}V(\mu) + C\max\{0, -\sum_{\theta\in\Theta}\lambda(\theta)\frac{\mu(\theta)}{\mu_0(\theta)}\})dp(\mu) && (18)\\ \text{s.t} \quad & p \in BP(\mu_0) \\ & V(\mu) \in \mathbb{V}(\mu) && \forall \mu \in \Delta(\Theta) \\ & \sum_{\theta\in\Theta}\lambda(\theta) = 1 \end{aligned}$$

Given any signaling scheme $p \in BP(\mu_0)$, since $\mathbb{V}(\mu)$ is a compact and convex set, and $\int_\mu(\sum_{\theta\in\Theta}\lambda(\theta)\frac{\mu(\theta)}{\mu_0(\theta)}V(\mu) + C\max\{0, -\sum_{\theta\in\Theta}\lambda(\theta)\frac{\mu(\theta)}{\mu_0(\theta)}\})dp(\mu)$ is convex in λ, and linear in $V(\mu)$, by Sion's minimax theorem we can interchange the $\max_V$ and $\min_\lambda$. Hence, we can obtain that $\mathcal{V}^*_C(\mu_0)$ can be solved by

$$\max_{p\in BP(\mu_0)} \min_{\lambda\in\text{aff}(\Theta)} \int_\mu \max\{\sum_{\theta\in\Theta}\lambda(\theta)\frac{\mu(\theta)}{\mu_0(\theta)}\max\mathbb{V}(\mu), \sum_{\theta\in\Theta}\lambda(\theta)\frac{\mu(\theta)}{\mu_0(\theta)}(\min\mathbb{V}(\mu) - C)\}dp(\mu). \tag{19}$$

This implies that we choose $V(\mu) = \max\mathbb{V}(\mu(m))$ if $\sum_{\theta\in\Theta}\lambda(\theta)\frac{\mu(\theta)}{\mu_0(\theta)} > 0$ and $V(\mu) = \min\mathbb{V}(\mu(m))$ if $\sum_{\theta\in\Theta}\lambda(\theta)\frac{\mu(\theta)}{\mu_0(\theta)} < 0$, which determines the best response selection rule. Thus, we have shown that

$$\mathcal{V}^*(\mu_0) = \max_{p\in BP(\mu_0)} \min_{\lambda\in\text{aff}(\Theta)} \mathcal{L}_C(\lambda, p).$$

The rest of the proof relies on Sion's minimax theorem as well. It is easy to verify that $BP(\mu_0)$ is a compact and convex set, and $\{\lambda|\lambda \in \text{aff}(\Theta)\}$ is a convex set. Moreover, $\mathcal{L}_C(\lambda, p)$ is linear in p and convex in λ since it is the maximum of two linear functions. Furthermore, $\mathcal{L}_C(\lambda, p)$ is continuous in λ. Since $\max\mathbb{V}(\mu)$ and $\min\mathbb{V}(\mu) - C$ are upper and lower semi-continuous, respectively, we have that $\sum_{\theta\in\Theta}\lambda(\theta)\frac{\mu(\theta)}{\mu_0(\theta)}\max\mathbb{V}(\mu)$ is upper semi-continuous when $\sum_{\theta\in\Theta}\lambda(\theta)\frac{\mu(\theta)}{\mu_0(\theta)} > 0$ and $\sum_{\theta\in\Theta}\lambda(\theta)\frac{\mu(\theta)}{\mu_0(\theta)}(\min\mathbb{V}(\mu) - C)$ is upper semi-continuous when $\sum_{\theta\in\Theta}\lambda(\theta)\frac{\mu(\theta)}{\mu_0(\theta)} < 0$. Therefore,

$$\hat{V}_{\lambda,C}(\mu) = \max\{\sum_{\theta\in\Theta}\lambda(\theta)\frac{\mu(\theta)}{\mu_0(\theta)}\max\mathbb{V}(\mu), \sum_{\theta\in\Theta}\lambda(\theta)\frac{\mu(\theta)}{\mu_0(\theta)}(\min\mathbb{V}(\mu) - C)\}$$

is upper semi-continuous and so is $\mathcal{L}_C(\lambda, p)$ in p. Hence, we can apply Sion's minimax theorem directly and complete the proof. □