A More Informed Sender Benefits the Receiver When the Sender Has Transparent Motives

Detailed Setup and Derivation

Mark Whitmeyer*

March 31, 2023

Model, Result, and Proof

$\theta \in \Theta$ denotes the (realized) state of the world, where Θ is some compact set. A sender (S) and a receiver (R) share a common prior, $\mu \in \Delta\Theta$, from which θ is drawn. S does not observe θ but instead privately observes the outcome of some statistical experiment about the state, $\pi\colon \Theta \to \Delta(Z)$, where Z is a set of signal realizations. S has a compact set of messages, M, available to him; and R has a compact set of actions, A, available to her. Both S and R have continuous utility functions. S has **Transparent Motives** (Lipnowski and Ravid (2020)); his utility depends only on the action taken by R: $u_S\colon A \to \mathbb{R}$. R's utility depends on her action and the state: $u_R\colon A \times \Theta \to \mathbb{R}$.

In the setting without mediation or delegation (i.e., the standard cheap-talk environment), S sends a message $m \in M$ to R, who subsequently chooses an action $a \in A$. Our solution concept in this setting is perfect Bayesian equilibrium, defined

*Arizona State University. Emails mark.whitmeyer@gmail.com

in the standard manner. In the setting with delegation, before S sends a message, R commits to a collection of conditional distributions over actions $\rho\colon M \to \Delta A$, which is observed by S, who subsequently sends a message. Our solution concept in this setting is subgame-perfect equilibrium obtained via backward induction. In the setting with mediation, before S sends a message, R commits to a signal, $\xi\colon M \to \Delta Y$, where Y is a set of signal realizations, which is observed by S, who subsequently sends a message. In this case, R does not observe S's message m but instead observes a signal realization z, which is drawn from $\xi(\cdot|m)$, before choosing an action $a \in A$. Our solution concept in this setting is subgame-perfect (Bayesian) equilibrium obtained via backward induction.

Let $\sigma(m|z)$ denote a mixed strategy chosen by S following signal realization z. Crucially, the collection of such mixed strategies following all $z \in \operatorname{supp}\pi$, $\sigma\colon Z \to \Delta M$, is itself an experiment. Given two experiments (stochastic maps) $\pi\colon \Theta \to \Delta Z$ and $\sigma\colon Z \to \Delta M$, we understand their composition, $\sigma \circ \pi\colon \Theta \to \Delta M$, to be the experiment

$$\sigma \circ \pi(m|\theta) = \sum_{z \in Z} \sigma(m|z)\pi(z|\theta).$$

Given experiment π and an equilibrium σ, the **Equilibrium Experiment** is $\sigma \circ \pi$.

We say that S is **More Informed** if he observes the outcome of statistical experiment π' instead of π, where π is a garbling of π'. That is π and π' are such that there exists a garbling η such that $\eta \circ \pi' = \pi$. We say that R is **Better Off** (as a result of a change in model primitives) if her maximal and minimal equilibrium payoffs weakly increase both when and when she cannot delegate or mediate.

Theorem 1. *If S is more informed, R is better off.*

Proof. In all three cases (delegation, mediation, and no delegation or mediation), regardless of S's information, there exists a babbling equilibrium in which S pools on some message following any $z \in \operatorname{supp}\pi$ and R chooses an action that is optimal under the prior with probability 1. This is the worst possible equilibrium for R, as no information is transmitted. Consequently, trivially, R's minimal equilibrium payoff weakly increases as S becomes more informed.

In the basic cheap-talk setting (no mediation or delegation), we can deduce a stronger result than the theorem:

Claim 2. *If τ is an equilibrium experiment given π, τ is an equilibrium experiment given π'.*

Proof. Fix π and some equilibrium σ that produces equilibrium experiment τ. As S has transparent motives, we must have

$$\mathbb{E}_{\rho(\cdot|m)}\left[u_S(a)\right] = \mathbb{E}_{\rho(\cdot|m')}\left[u_S(a)\right]$$

for all $m, m' \in \operatorname{supp}\sigma$, i.e., in the support of S's equilibrium strategy following some signal realization $z \in \operatorname{supp}\pi$; and where $\rho\colon M \to \Delta A$ is the collection of R's equilibrium behavioral strategies.

If S instead observes the realization of π', where $\pi = \eta \circ \pi'$, we construct an equilibrium as follows. We leave ρ unchanged and define a new collection of mixed strategies $\sigma' := \sigma \circ \eta$. These are incentive compatible for S as ρ is unchanged, and the set of messages sent on path is likewise unaltered. By construction, the equilibrium experiment is $\sigma' \circ \pi' = \sigma \circ \eta \circ \pi' = \tau$, as claimed. ■

The theorem (for the case of cheap talk) is a corollary of this claim: as any equilibrium experiment can be reproduced, R's equilibrium payoff must be the same as well.

Under delegation or mediation Claim 2 need not hold, as R's commitment strategy $\rho\colon M \to \Delta A$ (in the case of delegation) or R's chosen garbling of the message $\xi\colon M \to \Delta Y$ and subsequent behavioral strategy $\hat{\rho}\colon Y \to \Delta A$ (in the case of mediation), need no longer be part of an equilibrium when S is more informed, as ξ and ρ are chosen *ex ante*. However, as the receiver *could* always choose the ρ and ξ that are optimal when S's experiment is π, R's maximal equilibrium payoff must, therefore, be weakly higher when S's experiment is π', establishing the result. ■

ChatGPT-4 Proof Evaluation

Human: Can you evaluate a mathematical proof I wrote?

ChatGPT: Sure, I'd be happy to evaluate a mathematical proof you've written. Please provide me with the proof and any necessary context.

Human: Ok, do you know what a statistical experiment, also known as a Blackwell experiment or a stochastic map, is?

ChatGPT: Yes, I'm familiar with the concept of a statistical experiment or stochastic map. It refers to a mathematical model that describes a probabilistic process that generates a sequence of outcomes or observations. The process can be discrete or continuous and can involve a wide range of random variables and probability distributions. These models are often used in statistical inference, machine learning, and other areas of applied mathematics to simulate real-world phenomena or to make predictions based on incomplete or noisy data.

Human: Ok here is the setup for the proof.

ChatGPT: Great! Please provide the setup and the proof, and I'll do my best to evaluate it.

Human: There is a sender;

ChatGPT: Sure, please continue with the setup.

Human: and a receiver.

ChatGPT: I understand, there is a sender and a receiver. Please provide further context about the setup and the problem you are trying to solve.

Human: There is a state of the world, θ, which is drawn according to a commonly known prior from a compact set of possible states Θ.

ChatGPT: I see. So, there is a state of the world, denoted as θ, which is drawn from a compact set of possible states denoted as Θ. The state of the world is drawn according to a commonly known prior distribution. Please continue.

Human: The sender does not observe the state directly. Instead she privately observes the realization of a statistical experiment $\pi\colon \Theta \to \Delta(Z)$, where Z is a set of signal realizations and $\Delta(Z)$ is the probability simplex.

ChatGPT: I understand. So, the sender doesn't directly observe the state of the world, but instead privately observes the outcome of a statistical experiment denoted as π. This experiment maps each possible state of the world to a probability distribution over a set of signal realizations, denoted as Z, such that the probabilities for each signal realization add up to 1 (i.e., the probability simplex). Please continue with the setup.

Human: After observing a realization of this experiment, z, the sender sends a message to the receiver, m, where m is an element of a compact set M. Messages are cheap talk, they cost the sender nothing to send

ChatGPT: I understand. So, after observing the outcome of the statistical experiment π, denoted as z, the sender sends a message to the receiver. The message is denoted as m, and it belongs to a compact set of possible messages denoted as M. Importantly, sending a message costs the sender nothing, so messages are considered to be cheap talk. Please continue with the setup.

Human: The receiver observes the message, m, before taking an action a from a compact set of possible actions A. The receiver does not observe the state, nor the outcome of the statistical experiment.

ChatGPT: I understand. So, the receiver observes the message m sent by the sender, but doesn't observe the state of the world or the outcome of the statistical experiment π. The receiver then takes an action denoted as a, which belongs to a compact set of possible actions denoted as A. Please continue with the setup.

Human: The sender and receiver both have continuous utility functions. The sender's utility is a function only of the receiver's action. The receiver's utility depends both on the state and the receiver's action.

ChatGPT: I see. Both the sender and the receiver have continuous utility functions. The sender's utility function only depends on the receiver's action, denoted as a, and not on the state of the world or the outcome of the statistical experiment. The receiver's utility function, on the other hand, depends on both the state of the world, denoted as θ, and the receiver's action, denoted as a. Please continue with the setup.

Human: I say that if the sender observes the outcome of experiment π' rather than π, where π is a garbling of π', he is more informed. Alternatively, a more informed sender observes a more informative experiment. Do you understand this?

ChatGPT: Yes, I understand. You're saying that if the sender observes the outcome of a statistical experiment π', and π is a garbling of π', then the sender is more informed. In other words, a more informed sender is one who observes the outcome of a more informative statistical experiment. Please continue with the proof you'd like me to evaluate.

Human: I should note that this is a cheap talk game and the equilibrium concept is perfect Bayesian equilibrium. To remind you, the timing is that the sender first observes the outcome of a statistical experiment about an unknown state. The sender then sends a message to the receiver, who subsequently takes an action. Finally, the payoffs realize according to the specified utility functions.

ChatGPT: Thank you for the reminder. To summarize, we have a cheap talk game with perfect Bayesian equilibrium as the equilibrium concept. In this game, the sender first observes the outcome of a statistical experiment about an unknown state, then sends a message to the receiver, who subsequently takes an action. Finally, the payoffs are realized according to the specified utility functions. Please proceed with the proof you'd like me to evaluate.

Human: Here is the theorem that I will prove: the receiver's maximal payoff (so in the receiver-optimal perfect Bayesian equilibrium) does not decrease if the sender becomes more informed.

ChatGPT: Great! Please proceed with the proof.

Human: Fix a receiver-optimal equilibrium in the game when the sender's experiment is π.

ChatGPT: Okay, so we're fixing a receiver-optimal equilibrium in the game when the sender's experiment is π. Please continue.

Human: Observe that the sender's mixed strategy itself produces an experiment about the signal realizations the sender potentially observed. That is the sender's mixed strategy corresponds to an experiment $\sigma\colon Z \to \Delta M$, where recall Z is the set of signal realizations and M is the set of messages.

ChatGPT: I see. So, the sender's mixed strategy itself corresponds to an experiment about the signal realizations the sender potentially observed. This experiment is denoted as σ, and it maps each possible signal realization, denoted as z, to a probability distribution over the set of possible messages, denoted as M. Please continue with the proof.

Human: Thus, this equilibrium corresponds to a statistical experiment about the state of the world whose outcome is observed by the receiver, $\tau = \sigma \circ \pi \colon \Theta \to \Delta(M)$, where $\sigma \circ \pi$ is the experiment that is the composition of experiments π and σ.

ChatGPT: I understand. So, this equilibrium corresponds to a statistical experiment about the state of the world, denoted as τ, whose outcome is observed by the receiver. This experiment is obtained by composing the experiments π and σ, and is denoted as $\tau = \sigma \circ \pi$, where $\sigma \circ \pi$ is the experiment that maps each possible state of the world, denoted as θ, to a probability distribution over the set of possible messages, denoted as M. Please continue with the proof.

Human: Now suppose we make the sender more informed. Now she observes experiment π', which is such that π is a garbling.

ChatGPT: Understood. So, we're supposing that we make the sender more informed by having her observe experiment π', which is such that π is a garbling of π'. Please continue.

Human: Indeed, let $\pi = \xi \circ \pi'$ where ξ is a blackwell experiment and $\circ$ is the composition operator. Can you confirm that this is the definition that π is a garbling of π'?

ChatGPT: Yes, that's correct. When we say that π is a garbling of π', we mean that π is obtained by passing the signal realizations of π' through a blackwell experiment ξ. More precisely, we have $\pi = \xi \circ \pi'$, where $\circ$ denotes the composition of functions.

Human: Very good.

ChatGPT: Thank you! Please continue with the proof.

Human: Now, in the specified equilibrium when the sender observes experiment π, observe that the sender must be indifferent over every message in the support of the sender's strategy. Can you explain why that must be so?

ChatGPT: Yes, that's correct. In the specified equilibrium when the sender observes experiment π, the sender must be indifferent over every message in the support of the sender's strategy. This is because, in a perfect Bayesian equilibrium, the sender's mixed strategy must be such that it maximizes the sender's expected payoff given the receiver's optimal response. Since the sender's payoff function only depends on the receiver's action and not the sender's message, any two messages that result in the same optimal action for the receiver must yield the same expected payoff for the sender. Therefore, the sender must be indifferent over any two messages in the support of the sender's strategy that yield the same optimal action for the receiver.

Human: That is correct.

ChatGPT: Great! Please let me know how I can continue to assist you.

Human: Let me continue the proof.

ChatGPT: Sure, please proceed.

Human: Suppose when the experiment is π' we alter the sender's mixed strategy so that now her mixing produces experiment σ' where $\sigma' = \sigma \circ \xi$, where, recall, ξ is the experiment that garbles π' to produce π, i.e., $\pi = \xi \circ \pi'$.

ChatGPT: Okay, I follow. So now, with the experiment π' and the strategy σ', we have a new equilibrium. Can you explain why this new equilibrium must exist?

Human: Yes, I can. First, let us assume that the receiver's strategy is unchanged (we will confirm that this works shortly). Second, the sender is randomizing over the same set of messages, which the sender is still willing to do, because the receiver's behavior is unchanged, and we already have stated that the sender is indifferent about sending these messages. Clear?

ChatGPT: Yes, that's clear.

Human: Finally, observe the sender's strategy begets a statistical experiment about the state of the world. Recall that when the sender observed experiment π, this statistical experiment was $\sigma \circ \pi$. Now that the sender observes experiment π' and randomizes in the new specified way, this statistical experiment is $\sigma' \circ \pi'$. However, this is the same statistical experiment as $\sigma \circ \pi$ due to our specification that $\sigma' = \sigma \circ \xi$, i.e., $\sigma' \circ \pi' = \sigma \circ \xi \circ \pi' = \sigma \circ \pi$. As the receiver is observing the outcome of the same experiment as previously, the receiver is still willing to respond in the same way to each message sent by the sender.

ChatGPT: Yes, that is correct. So the receiver's behavior is unchanged under the new equilibrium. Therefore, the only change is in the sender's strategy, which leads to the same experiment being observed by the receiver. Hence, the receiver's maximal payoff does not decrease, and the proof is complete.

Human: Woohoo!

ChatGPT: Congratulations on the successful proof! Is there anything else you would like to discuss or explore?

Human: No, thank you!

ChatGPT: You're welcome! It was a pleasure helping you.

References

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