A Rothschild-Stiglitz approach to Bayesian persuasion

Matthew Gentzkow and Emir Kamenica*
Stanford University and University of Chicago

September 2015

Abstract

Rothschild and Stiglitz (1970) introduce a way to represent random variables as convex functions (integrals of the cumulative distribution function). Combining their result with Blackwell's Theorem (1953), we characterize the set of distributions of posterior means that can be induced by a signal. This characterization provides a novel way to analyze a class of Bayesian persuasion problems.

JEL classification: D83
Keywords: mean-preserving spreads; information design

*We thank the Sloan Foundation and the University of Chicago Booth School of Business for financial support.

1 Introduction

Consider a situation where one person, call him Sender, generates information in order to persuade another person, call her Receiver, to change her action. Sender and Receiver share a common prior about the state of the world. Sender can publicly generate any signal about the state and Receiver observes the signal realization before she takes her action.[1]

Kamenica and Gentzkow (2011) analyze a general version of this 'Bayesian persuasion' problem.[2] They draw on an insight from Aumann and Maschler (1995) to develop a geometric approach to Sender's optimization problem. They derive a value function over beliefs and then construct the optimal signal from the concavification of that value function.[3] This approach provides ample intuition about the structure of the optimal signal, but has limited applicability when the state space is large. The dimensionality of the space of beliefs is roughly the same as the cardinality of the state space,[4] so the value function and its concavification can be visualized easily only when there are two or three states of the world. When the state space is infinite, the concavification approach requires working in an infinite-dimensional space.

In this paper we analyze a class of Bayesian persuasion problems where the state space may be large but Sender and Receiver's preferences take a simple form: the state ω is a random variable, Receiver's optimal action (taken from a finite set) depends only on $\mathbb{E}[\omega]$, and Sender's preferences over Receiver's action are independent of the state.

This environment captures a number of economically relevant settings. For example, it might be the case that Sender is a firm, Receiver is a consumer, and ω is the match quality between the attributes of firm's product and the consumer's preferences. Since the latter are unknown to the firm, the common prior assumption is palatable. The interpretation of the signal in this case is the firm's choice of what information about the product to provide to the consumer. For example, a software company can decide how many features to provide in the trial version of its product.

Kamenica and Gentzkow (2011) also examine this environment, but do not characterize the

[1] A *signal*, in our terminology, is a map from the true state of the world to a distribution over some signal realization space. Others terms for a signal include *experiment* and *signal structure*.

[2] Gentzkow and Kamenica (2014) extend the analysis to the case of costly signals and Gentzkow and Kamenica (2015) to a situation with multiple senders.

[3] A *concavification* of a function f is the smallest concave function that is everywhere weakly greater than f.

[4] When the state space Ω is finite, the space of beliefs has $|\Omega| - 1$ dimensions.

optimal signal. They show that if one considers the value function over Receiver's posterior mean, the concavification of that value function pins down whether Sender can benefit from generating information but does not determine the optimal signal.

The problem is that it is difficult to characterize the set of feasible distributions of the posterior mean. Any distribution of posterior beliefs whose expectation is the prior can be induced by some signal; but, it is not possible to induce every distribution of posterior means whose expectation is the prior mean.

In this paper, we combine insights from Blackwell (1953) and Rothschild and Stiglitz (1970) to derive the characterization of all feasible distributions of the posterior mean. We then use this characterization to analyze the aforementioned class of Bayesian persuasion problems.

Kolotilin (2014) and Kolotilin *et al.* (2015) examine closely related environments. They make the same assumptions on preferences but allow for Receiver to have private information. They focus exclusively on the case with a binary action. Kolotilin (2014) shows that neither Sender's nor Receiver's payoff is necessarily monotone in the precision of Receiver's private information. Kolotilin *et al.* (2015) consider "private persuasion" where Receiver reports his private type before Sender generates information. They show that Sender never strictly benefits by allowing for private persuasion.[5] While the focus of these papers is somewhat different, our proof draws on a result in Kolotilin (2014).

2 The model

The state of nature is a random variable ω on $[0,1]$. Sender and Receiver share a common prior F_0 on ω. Throughout the paper we denote any distribution over real numbers by its cumulative distribution function (CDF); hence, under the prior $Pr(\omega \leq x) = F_0(x)$. Let m_0 denote the mean of F_0. A *signal* π consists of a *signal realization space* S and a family of distributions $\{\pi_\omega\}$ over S. Sender chooses a signal. Receiver observes the choice of the signal π and the signal realization s. Receiver then chooses an action from a finite set. Her optimal action depends on her expectation of the state, $\mathbb{E}[\omega]$. Without loss of generality we label the actions so that action a_i is optimal

[5] Gupta (2014) and Wang (2015) also contrast public and private persuasion but they have a different notion of private persuasion. Specifically, they consider a case with multiple receivers and contrast the case where all of them observe the same signal realization with the case where they observe independent draws of the signal.

Figure 1: Signals as convex functions

(a) CDFs of posterior means (b) Functions induced by signals

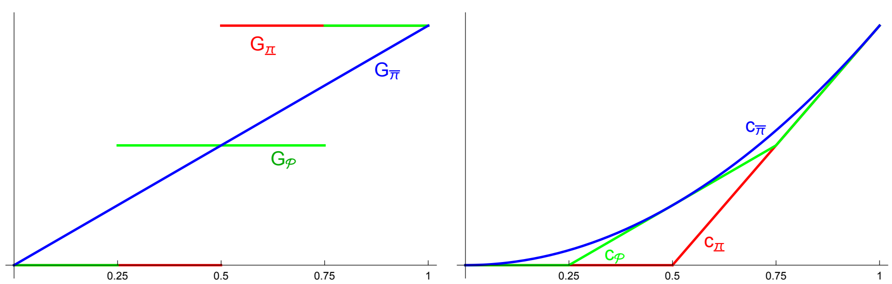

if $\gamma_i \leq \mathbb{E}[\omega] \leq \gamma_{i+1}$ given some set of cutoffs $\gamma_0 \leq \gamma_1 \leq \ldots \leq \gamma_n \in [0,1]$. Sender has some state-independent utility function over Receiver's action.

3 Signals as convex functions

Given a signal π, a signal realization s induces a posterior F_s. For each signal realization, let m_s denote the mean of F_s. A signal induces a distribution of posteriors and hence a distribution of posterior means. Let $G_\pi : \mathbb{R} \to [0,1]$ denote the distribution of posterior means induced by signal π. Then, for each signal π, let c_π denote the integral of G_π, i.e., $c_\pi(x) = \int_0^x G_\pi(t)\,dt$. If c_π is thus obtained from π we say that π *induces* c_π.

We illustrate this definition with some examples. Suppose that F_0 is uniform. Consider a totally uninformative signal $\underline{\pi}$. This signal induces a degenerate distribution of posterior means always equal to $m_0 = \frac{1}{2}$. Hence, $G_{\underline{\pi}}$ is a step function equal to 0 below $\frac{1}{2}$ and equal to 1 above $\frac{1}{2}$. The convex function $c_{\underline{\pi}}$ induced in turn is thus flat on $\left[0, \frac{1}{2}\right]$ and then linearly increasing from $\frac{1}{2}$ to 1. At the other extreme, consider a fully informative signal $\overline{\pi}$ that fully reveals the state. In that case, each posterior has a degenerate distribution with all the mass on the true state and thus $G_{\overline{\pi}} = F_0$. Since F_0 is uniform, $G_{\overline{\pi}}$ is linear, and thus $c_{\overline{\pi}}$ is quadratic: $c_{\overline{\pi}}(x) = \frac{1}{2}x^2$. Finally, consider a "partitional" signal $\mathscr{P}$ that gives a distinct signal realization depending on whether the state is in $\left[0, \frac{1}{2}\right]$, or $\left(\frac{1}{4=2}, 1\right]$. Then, $G_{\mathscr{P}}$ is a step function and $c_{\mathscr{P}}$ is piecewise-linear. Figure 1 depicts these CDFs and functions.

If we consider an arbitrary signal π, what can we say about c_π? Since G_π is a CDF and thus increasing, c_π as its integral must be convex. Moreover, since any signal π is a garbling of $\overline{\pi}$, we must have that $G_{\overline{\pi}}$ is a mean-preserving spread of G_{π} (Blackwell 1953); hence, $c_{\overline{\pi}} \geq c_{\pi}$ by Rothschild and Stiglitz (1970). Similarly, since $\underline{\pi}$ is a garbling of π, G_{π} is a mean-preserving spread of $G_{\overline{\pi}}$ and thus $c_{\pi} \geq c_{\underline{\pi}}$. Putting these observations together, we obtain:

Remark 1. Given any signal π, the induced function c_{π} is convex. Moreover, $c_{\overline{\pi}}(x) \geq c_{\pi}(x) \geq c_{\underline{\pi}}(x)$ $\forall x \in [0,1]$.

Note that this result applies for any prior, not just for the uniform example depicted in Figure 1. In general, functions $c_{\underline{\pi}}$ and $c_{\overline{\pi}}$ are determined by the prior with $c_{\underline{\pi}}$ flat to the left of m_0 and then increasing with a slope of 1, and $c_{\overline{\pi}}$ equal to the integral of F_0.

4 Convex functions as signals

Now suppose that we are given some function that satisfies the properties from Remark 1. Is it always the case that there is some signal that induces this function? The answer to this question turns out to be affirmative:

Proposition 1. *Given any convex function* $c : [0,1] \to \mathbb{R}$ *such that* $c_{\overline{\pi}}(x) \geq c(x) \geq c_{\underline{\pi}}(x)$ $\forall x \in [0,1]$*, there exists a signal that induces it.*

Proof. Consider some function c satisfying the given properties. Define a function

$$G(x) = \begin{cases} 0 & : \quad x < 0 \\ c'(x) & : \quad 0 \leq x < 1 \\ 1 & : \quad x \geq 1 \end{cases}$$

where $c'(x)$ denotes the right derivative of c at x. Since c is convex, its right derivative must exist. Moreover, since c is convex and $0 \leq c'(x) \leq 1$ for all x (cf: Lemma 1 in the Appendix), G is increasing and right-continuous. We also clearly have that $lim_{x\to-\infty}G(x) = 0$ and $lim_{x\to\infty}G(x) = 1$. Hence, G is a CDF. Now, since $\int_0^x F_0(t)\,dt = c_{\overline{\pi}}(x) \geq c(x) = \int_0^x G(t)\,dt$, we know that F_0 is a mean-preserving spread of G. By Proposition 1 in Kolotilin (2014), this in turn implies that there must exist a signal that induces G as the distribution of posterior means. □

Proposition 1 thus provides us with a simple characterization of the distributions of posterior

Figure 2: Feasible distributions of posterior means vs. all random variables with $\mathbb{E}[\tilde{m}] = m_0$

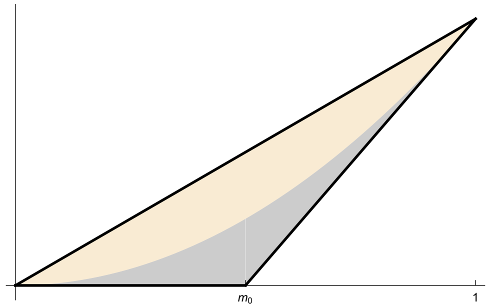

means that can be induced by a signal. Figure 2 contrasts the space of functions induced by all random variables whose expectation is the prior mean (any convex function in the lightly shaded area) with the subset of those that represent feasible distributions of the posterior means (any such function in the darker area in the bottom right).

5 Optimal signals

In the previous section we transformed Sender's "budget set" of all signals into a "budget set" of convex functions. Our next step is to analyze how to determine Sender's payoff for a given function in this new budget set.

The key observation is that – under the preference structure we have assumed – Sender's payoff is entirely determined by the local property of the induced function around the action cutoffs. In particular, the left derivative of c_π at γ_i fully determines how often Receiver takes an action in the set $\{a_i, a_{i+1}, ..., a_n\}$. Hence, once we know $c'_\pi(\gamma_i)$ for each i – where, from here on, we let c'_π denote the left derivative – we can back out Sender's payoff.

5.1 Two actions

Consider the simplest case where Receiver takes one of two actions: a_0 or a_1. To make the problem non-trivial, we assume that Sender prefers a_1, but $m_0 < \gamma_1$.[6] In this case, Sender wants to design a signal that maximizes the probability of a signal realization s such that $\mathbb{E}_{F_s}[\omega] \geq \gamma_1$.[7] This simple problem can also be solved algebraically (Ivanov 2015).[8] We nonetheless begin with this simplest

[6]Otherwise, a completely uninformative signal is clearly optimal.

[7]It is easy to show that in any equilibrium, Receiver must break her indifference at γ_1 in Sender's favor.

[8]An optimal signal is a partition that reveals whether ω belongs to $[x^*, 1]$, with x^* defined by $\int_{x^*}^1 x dF_0(x) = \gamma_1$.

Figure 3: Optimal signals with binary actions

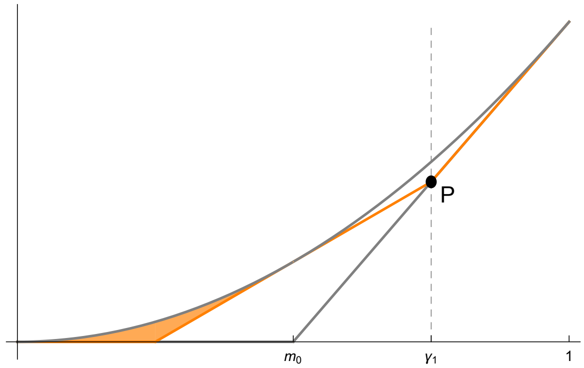

example as it illustrates our approach in the most transparent way.

As mentioned above, if Sender induces c_π with $c'_\pi(\gamma_1) = k$, his payoff will be proportional to $1 - k$. Hence, Sender wants to induce a function that minimizes the left derivative at γ_1. As we can see Figure 3, he cannot bring this derivative all the way down to zero. Doing so would violate the restriction that c_π must be convex, and bounded above by $c_{\overline{\pi}}$. In fact, looking at Figure 2, it is easy to see that any optimal c_π – the one that minimizes the left derivative – must satisfy two features. First, it must coincide with $c_{\underline{\pi}}$ at γ_1, as indicated by the "pivot point" labeled P in Figure 3. Second, the "arm" leaving P to the left should be "pivoted up" as much as possible, until it is tangent to $c_{\overline{\pi}}$. This identifies the optimal signals since the behavior of the function to the left of the tangency point is irrelevant. Thus, any convex function within the shaded area of Figure 3 is optimal. These functions correspond to signals that yield a single, deterministic realization s when ω is above the tangency point and generate arbitrary (potentially stochastic) realizations (not equal to s) for other states. The top of the shaded area is induced by a signal that fully reveals all ω below the tangency point while the bottom of the area is induced by a signal that always generates a single realization for all those states.

5.2 More actions

Now suppose Receiver can take one of three actions and Sender's utility is 0 from a_0, 1 from a_1, and $\lambda > 1$ from a_2. Suppose that $m_0 \in (\gamma_1, \gamma_2)$. Looking at Figure 4, we first note that the optimal function must go through point P. The only question that remains, therefore, is where the function should cross γ_1 – this point determines the tradeoff between how frequently a_1 and a_2 are taken.

Figure 4: Optimal signal with three actions

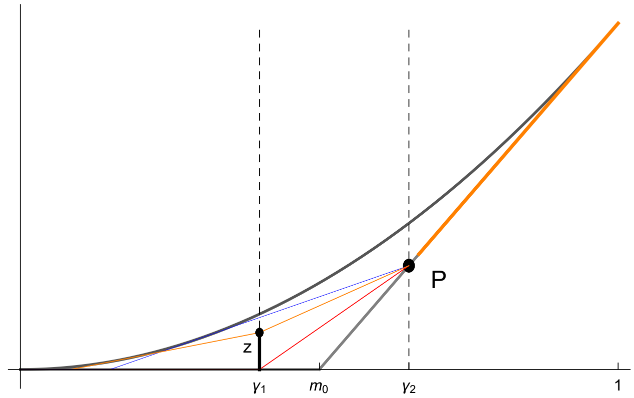

At one extreme, we have the blue function that maximizes the probability of a_2. This occurs at the expense of a_1 never happening. At the other extreme is the red function that ensures that a_0 never happens, but consequently leads to a_2 being less frequent than it could be. Finally, the orange function shows a "compromise" solution where all three actions are taken with positive probability. As the Figure shows, we can index all potentially optimal functions with a single-dimensional parameter z that denotes the height at which the function crosses γ_1.

How does Sender's payoff vary with z? Probability of a_2 is one minus the slope of the second segment, which is linearly increasing in z.[9] Probability of a_1, on the other hand, is decreasing in z. This relationship is generally not linear. As can be seen from Figure, it depends on F_0. In the Appendix, we explicitly compute the relationship between z and the probability of a_1 for the case of a uniform prior. It takes the form of $A - \sqrt{B - 2z} - Cz$ where A, B, and C are constants that depend on γ_1 and γ_2. Because the relationship is not linear, we do not necessarily end up at a corner solution with either the blue line (zero probability of a_1) or the red line (zero probability of a_0) being optimal.[10] For example, if the prior is uniform, $\gamma_1 = \frac{1}{3}$, $\gamma_2 = \frac{2}{3}$, and $\lambda = 3$, the optimal z is $\frac{1}{24}$. This function cannot be induced through an interval partition. One signal that that achieves the optimum is a non-monotone partition that reveals whether the state is in $\left[0, \frac{8}{48}\right]$ inducing a_0, in $\left[\frac{11}{48}, \frac{21}{48}\right]$ inducing a_1, or in $\left[\frac{8}{48}, \frac{11}{48}\right] \cup \left[\frac{21}{48}, 1\right]$ inducing a_2.

[9] Specifically, the probability of a_2 is $1 - \frac{\gamma_2 - m_0 - z}{\gamma_2 - \gamma_1}$.

[10] Of course if λ is particularly high or particularly low, a corner solution can be optimal. Also, we note that Sender's payoff is linear in the induced function, so the optimum is always an extreme point of the feasible set.

6 Conclusion

Previous work on Bayesian persuasion built on the observation that a distribution of posterior beliefs is feasible, i.e., can be induced by a signal, if and only if its expectation is the prior. In this paper, we characterize the set of feasible distributions of posterior means. This provides us with a novel way to solve an important class of Bayesian persuasion problems.

7 Appendix

7.1 Additional proofs

Lemma 1. *Consider any convex function* $c : [0,1] \to \mathbb{R}$ *such that* $c_{\overline{\pi}}(x) \geq c(x) \geq c_{\underline{\pi}}(x)$ $\forall x \in [0,1]$. *Let* $c'(x)$ *denote the right derivative of* c *at* x. *Then,* $0 \leq c'(x) \leq 1$ $\forall x \in [0,1]$.

Proof. First note that since c is convex, it must be continuous.

Suppose that $c'(x^*) < 0$ for some $x^* \in [0,1]$. Since c is convex, this implies that $c'(0) < 0$. Since $c_{\overline{\pi}} \geq c$ and $c_{\overline{\pi}}(0) = 0$, we have $c(0) \leq 0$. Thus, since $c'(0) < 0$, for a small enough x, we have $c(x) < 0$. But since $c_{\underline{\pi}} \geq 0$, this violates the assumption that $c \geq c_{\underline{\pi}}$.

Suppose that $c'(x^*) > 1$ for some x^*. Since c is convex, we have that $c'(1) > 1$. Since $c_{\overline{\pi}} \geq c \geq c_{\underline{\pi}}$.and $c_{\overline{\pi}}(1) = c_{\underline{\pi}}(1) = 1 - m_0$, we have $c(1) = c_{\underline{\pi}}(1)$. Now we consider two cases. First, suppose that $m_0 < 1$. In that case, $c'_{\underline{\pi}}(1) = 1$. But then $c'(1) > 1$ and $c(1) = c_{\underline{\pi}}(1)$ jointly imply that there exists an $x < 1$ s.t. $c < c_{\underline{\pi}}$ so we have reached a contradiction. Alternatively, suppose that $m_0 = 1$. In that degenerate case, we have that $c_{\overline{\pi}}(x) = c_{\underline{\pi}}(x) = 0$ for all x, so we must have $c(x) = 0$ for all x. □

7.2 Optimal signal with three actions

Suppose F_0 is uniform on $[0,1]$. Consider a function $f(x) = a + bx$ that is tangent to $c_{\overline{\pi}}$ to the left of γ_1 and crosses through the point (γ_1, z). Since F_0 is uniform, we know $c'_{\overline{\pi}}(x) = x$. Hence, f must be tangent to $c_{\overline{\pi}}$ at b and we have $f(b) = c_{\overline{\pi}}(b)$, which means $f(x) = -\frac{b^2}{2} + bx$. Since $f(\gamma_1) = z$, we have $\frac{-b^2}{2} + b\gamma_1 = z$. By the quadratic equation, this implies $b = \gamma_1 \pm \sqrt{\gamma_1^2 - 2z}$. From Figure 4 we can clearly see that b is increasing in z, so we know that that the smaller solution is the correct one: $b = \gamma_1 - \sqrt{\gamma_1^2 - 2z}$.

Since the probability of a_0 is b and the the probability of a_2 is $1 - \frac{\gamma_2 - \frac{1}{2} - z}{\gamma_2 - \gamma_1}$, the probability of a_1 is $\frac{\gamma_2 - \frac{1}{2} - z}{\gamma_2 - \gamma_1} - \gamma_1 + \sqrt{\gamma_1^2 - 2z}$. If we have $\gamma_1 = \frac{1}{3}$, $\gamma_2 = \frac{2}{3}$, and $\lambda = 3$, the overall payoff is $\left(\frac{\gamma_2 - \frac{1}{2} - z}{\gamma_2 - \gamma_1} - \gamma_1 + \sqrt{\gamma_1^2 - 2z}\right) + \left(1 - \frac{\gamma_2 - \frac{1}{2} - z}{\gamma_2 - \gamma_1}\right) 3$, which is maximized at $z = \frac{1}{24}$.

References

Aumann, Robert J., and Michael B. Maschler. 1995. *Repeated Games with Incomplete Information.* Cambridge, MA: MIT Press

Blackwell, David. 1953. Equivalent comparisons of experiments. *The Annals of Mathematical Statistics.* Vol. 24, No. 2, 265-272.

Gentzkow, Matthew and Emir Kamenica. 2014. Costly persuasion. *American Economic Review: Papers & Proceedings.* Vol. 104, No. 5, 457-462.

Gentzkow, Matthew and Emir Kamenica. 2015. Competition in persuasion. Working paper.

Gupta, Seher. 2014. Information design in collective decision games. Working paper.

Ivanov, Maxim. 2015. Optimal signals in Bayesian persuasion mechanisms with ranking. Working paper.

Kamenica, Emir and Matthew Gentzkow. 2011. Bayesian persuasion. *American Economic Review.* Vol. 101, No. 6, 2590-2615.

Kolotilin, Anton. Optimal information disclosure: Quantity vs. quality. Working paper.

Kolotilin, Anton, Ming Li, Tymofiy Mylovanov, and Andriy Zapechelnyuk. 2015. Persuasion of a privately informed receiver. Working paper.

Rothschild, Michael and Joseph Stiglitz. 1970. Increasing Risk: I. A definition. *Journal of Economic Theory.* Vol 2, 225-243.

Wang, Yun. 2015. Bayesian persuasion with multiple receivers. Working paper.