Persuasion with Multiple Actions*

Davit Khantadze† Ilan Kremer‡ Andrzej Skrzypacz§

June 25, 2021

Abstract

We consider a Bayesian persuasion model in which the receiver takes multiple actions. We compare a simultaneous procedure, in which the receiver takes all actions after the realization of a single signal, to a sequential procedure in which the receiver receives information gradually and takes the actions sequentially. We characterize conditions under which the sequential procedure leads to a higher payoff and when it achieves a maximal outcome despite potential information leakage. We also discuss the optimal ordering of actions.

JEL Codes: D21

Keywords: .

1 Introduction

In this paper, we study optimal persuasion when the receiver takes multiple actions (or multiple receivers take actions). We characterize the optimal

*We thank seminar participants, our colleagues, and students for comments and suggestions.

†Email: d.khantadze@gmail.com.

‡The Hebrew University and University of Warwick. Email: ikremer@huji.ac.il.

§Stanford Graduate School of Business, Stanford University. Email: skrz@stanford.edu

dynamic persuasion and show when and why gradual/sequential persuasion improves upon one-time communication. Our model and results apply to a wide range of situations, from a salesperson persuading a buyer to buy multiple products to a company manager persuading various constituencies (e.g., a regulator and investors).

Consider the classic example of Bayesian persuasion from the introduction to Kamenica and Gentzkow (2011). A judge, who is the 'receiver,' finds a defendant guilty if and only if his belief that the defendant committed a crime is above a certain threshold. The prosecutor, who is a 'sender,' designs an objective test, and the judge will observe the realization of that test. The sender's objective is to maximize the probability of conviction. In our model, the defendant is accused of (at least) two crimes, A and B. The judge decides on both offenses. The prosecutor wishes to maximize the expected number of convictions. We show how the multiplicity of actions affects optimal persuasion and the benefits of gradual persuasion.

Our goal is to compare sequential with simultaneous procedures. In a simultaneous procedure, the sender provides a single signal, and the receiver decides simultaneously about both A and B. A sequential procedure has two sequential steps. In the first step, the sender provides one signal, and the receiver decides on A. In the second step, the prosecutor reveals the realization of a second signal, and the receiver decides on B. We assume that the receiver cannot reverse an action he has taken. If he could reverse his decisions, there would not be a difference between sequential and simultaneous procedures.

While a simultaneous procedure can be captured by the classic Bayesian persuasion framework with a sufficiently rich decision/signal space, a sequential procedure cannot. When the priors regarding different actions are not independent, there are feedback effects. Information on whether the defendant has committed the first crime informs the judge about the possibility that the defendant has committed the second crime. Similarly, the willingness to buy good A is correlated with the willingness to buy good B. Such correlation affects the way the receiver behaves both in the simultaneous and sequential procedure.

As we show, the sender weakly prefers a sequential procedure to a simultaneous one. Hence, our goal is to examine when and to what extent a sequential procedure strictly outperforms the simultaneous one. We show that the sequential procedure is strictly improving if and only if the optimal simultaneous procedure fails to achieve a certain upper bound on sender payoffs. It is constructed by considering the minimal information leakage from the persuasion that would be achieved if prior beliefs about the different states were independent. We refer to that upper bound as the first best payoff, the best payoff the sender could have received if he faced no information leakage constraints.

In the first part of the paper, we assume that the sequential procedure is based on the order of actions predetermined by the sender. We later relax this assumption and show that if the sender can choose the order of the receiver's decisions in reaction to the signals' realizations, then the sender can always achieve the first-best payoff. In particular, we construct a 'Pacman Procedure' that enables the sender to achieve the first-best payoff via a sequence of binary signals.

We finish by examining when a simple sequential procedure could achieve the first-best payoff, and if so, what is the optimal order of persuasion.

To gain intuition about our model and results, we start with three examples.

1.1 Examples

In the following three examples, we consider a seller who wishes to sell two goods, A and B. The goods are either of high or low quality. The seller has zero value for the goods (regardless of quality). The buyer has zero value for a low-quality good and a positive value for a high-quality good. Prices are fixed and normalized to one. The buyer follows a threshold rule where he buys if his belief of high quality exceeds a certain threshold (so that the expected value of the good is above the price). The seller's payoff equals the expected number of goods he sells.

Example 1. *The buyer values a high-quality good at* 1.25*. So, he buys a good if and only if his belief of high quality exceeds* $1/1.25 = 0.8$*. The prior is that exactly one of the two goods is of high quality and the other is low quality. The probability that each of the two goods is of high quality is* 0.5*.*

Consider first the simultaneous procedure. If the seller reveals no information, the buyer does not buy any good, and the seller's payoff is zero. If instead, the seller provides full information, he would sell one of the goods, and his payoff would be one. This is the optimal simultaneous procedure as the buyer will never buy both goods. If he buys good A*, then he believes that with a probability of at least* 0.8 *it is of high quality. This implies that the probability that* B *is of high quality does not exceed* 0.2 *and he does not buy it.*

Consider now a sequential procedure. As mentioned above, we assume that a transaction cannot be reversed. The seller first provides a signal about A*. This signal induces a posterior belief that* A *has high quality of either* 0 *or* 0.8 *with probabilities* $3/8$ *and* $5/8$*, respectively. Conditional on the posterior of* 0*, the buyer would not buy* A *in the first step but buy* B *in the second step. Conditional on the posterior of* 0.8*, the buyer buys* A *in the first step, and he assigns a probability of* 0.2 *that* B*'s quality is high. In the second step, the seller can provide more information that, with some probability, will convince the buyer also to buy* B*. In particular, given a prior of* 0.2*, he can induce a posterior of* 0.8 *with probability* $1/4$*. This sequential procedure induces the buyer to always buy at least one of the goods and, in some cases buy both goods. The expected number of goods sold is* $3/8 + 5/8 * (1 + 1/4) = 1.156$*.*

Note that in this example, the upper bound for the sum of seller and buyer payoffs is 1.25 (this is the total surplus if the trade is efficient). So 1.25 is also an upper bound on the seller payoffs. In the optimal simultaneous procedure, the seller's payoff is 1, and the buyer's surplus is 0.25. In the sequential procedure we described, the seller's payoff is 1.156 while the buyer's surplus is 0.094 (the sum is still 1.25). While the seller does not extract the entire surplus, we show in the appendix that it is the optimal sequential procedure for the seller.

Example 2. *Similar to example 1, but the value of high-quality good A is* $5/3$*, and high-quality good B is* $5/4$*. Hence, the thresholds for goods A and B are* 0.6 *and* 0.8*, respectively. The prior is that the qualities are perfectly correlated: good A has high quality if and only if good B does. The prior that good A (and B) is of high quality is* 0.5*.*

Consider first the simultaneous procedure. Given the perfect correlation, the number of goods sold is a function of the probability that the common quality is high. If it is lower than 0.6*, then no goods will be sold. If it is between* 0.6 *and* 0.8*, one good (good A) will be sold. If it exceeds* 0.8*, both goods will be sold. Given the prior of* 0.5*, using the standard concavification argument we conclude that the optimal procedure is to generate a posterior of* 0.8 *with probability* $5/8$ *and zero with a probability of* $3/8$*. The expected number of goods sold is then* $2 * 5/8 = 1.25$*.*

Consider now a sequential procedure in which the seller tries first to sell good A. In the first step, the seller induces a posterior of 0.6 *with probability* $5/6$ *and a posterior of* 0 *with probability* $1/6$*. He thus sells good A with a probability of* $5/6$*. In the second step, if he induced a zero posterior, he will not sell good B. If he has induced a prior of* 0.6 *then he can induce a final posterior of* 0.8 *with a probability of* $6/8 = 3/4$ *(and a posterior of* 0 *otherwise). Hence, good B is sold with probability* $(3/4) * (5/6) = 5/8$*. The expected total sales are* $5/6 + 5/8 = 1.458$*. Note that this equals the total available surplus, which is an upper bound for the seller's payoff.*

Example 3. *Same as in example 2, but the qualities of the two goods are independent. That is, each good is of high quality with a probability of* 0.5*, but the quality of A is independent of B. Consider a simultaneous procedure in* *which the seller induces a posterior for A of* 0.6 *with probability* 5/6 *and zero with probability* 1/6*. Given the independence, the seller can simultaneously induce a posterior for B of* 0.8 *with probability* 5/8 *and zero with probability* 3/8*. The buyer then buys A with a probability* 5/6 *and B with a probability* 3/4*, for total expected sales of* 1.458*. As the seller extracts the total available surplus, no sequential procedure can improve upon this.*

Examples 1 and 2 highlight a key reason why a sequential procedure generates a higher payoff for the sender: the buyer sometimes regrets his decision in the first step. A key difference between the two examples is that in the second example, the seller extracts the entire surplus in the sequential procedure. In the first example, he does not.

The seller in example 3 extracts the entire surplus with both procedures. The role of that example is to illustrate that the seller's payoff in any problem is bounded from above by a payoff in a game with a modified prior. Under that modified prior, the states are independent, and the marginal distributions are the same as in the original prior. We show formally (with more decisions and arbitrary payoffs) that in the case of the modified prior sender's payoffs are the same for the optimal sequential and simultaneous procedures. Moreover, these payoffs bound from above the sender's payoffs from any of the two procedures for the original prior. The intuition is that, when beliefs are independent, persuading the receiver on one of the dimensions does not change the receiver's beliefs about other dimensions. In contrast, with a correlated prior, revealing information about one of the dimensions reveals information about another, imposing additional constraints on the seller's ability to persuade. We refer to this effect as the information leakage problem.

In Section 3 we prove the main result of our paper, that the optimal sequential procedure strictly dominates the optimal simultaneous one if and only if the simultaneous procedure fails to achieve the first-best payoff (the upper bound). The following table illustrates this result:

	Simultaneous	Sequential	First-Best
Example 1	1	1.156	1.25
Example 2	1.25	1.458	1.458
Example 3	1.458	1.458	1.458

It is important to note that (i) this result is not based on the optimal sequential procedure always achieving the first-best payoffs, as can be seen in Example 1, and (ii) we prove this result for a general action space so this extends beyond a buyer/seller interaction with binary decisions.

In the main model, we assume that in the sequential procedure the sender first chooses the sequence of decisions and then releases (sequentially) information. That is, in the examples above, the seller first chooses the sequence of goods the buyer has to decide about and only then releases the information. In Section 4 we allow even more complicated dynamic persuasion: the sender is allowed to send many signals and choose the order of actions based on the realizations of these signals. The main result in that section is that this more flexible procedure always allows the sender to achieve the first-best payoff, no matter the correlation between the states.

Finally, in Section 5 we discuss when the simple sequential procedure can and cannot achieve the first-best payoff, and that also allows us to provide some intuition about the optimal sequencing of actions.

1.2 Literature Review

Our paper contributes to the literature on information design and Bayesian persuasion. Aumann et al. [1995], Kamenica and Gentzkow [2011], and Bergemann and Morris [2016a,b] are some of the first models in this field. We apply this methodology to settings where the decision maker takes multiple irreversible actions and information can be provided sequentially.

By considering sequential persuasion of multiple actions, we contribute to the literature on dynamic persuasion, such as Hörner and Skrzypacz [2016] and Ely [2017]. In that area, our work is particularly related to Ely and Szydlowski [2020] and Orlov et al. [2020].[1] These papers analyze the provision of information in a dynamic setting where the decision maker's actions are irreversible. In those papers, the receiver decides in every period whether to stop or continue (stop putting effort or take an irreversible decision to remove a product from a market). The decisions to stop or continue are not reversible because the receiver cannot get back in time. There are three main differences between our model and those two models. Most importantly, in our model, the set of actions available to the receiver in later steps does not depend on the actions taken previously. In contrast, in those papers, if the receiver stops in time t, he has no other choices to make at later times. Second, we do not allow the sender to condition the past signals on the previous actions taken by the receiver. In that sense, we do not allow information to be used as a carrot. In Ely and Szydlowski [2020] and Orlov et al. [2020], a major part of the intuition why gradual information release helps the sender is precisely the carrot aspect of future information: the sender can entice the receiver to delay stopping by promising additional information in the future. In contrast, in our model, the receiver takes myopically optimal action for every decision problem (because future payoffs do not depend on the current action). In this way, we add to the intuitions in those two papers: sequential information disclosure can help the sender beyond creating a promise of future information.[2]

Our results are related to Bergemann and Morris [2016b] who show that if the receiver has access to additional information, the set of feasible outcomes is reduced. The connection to our paper is that since the sender provides information about multiple correlated states, information released about state

[1]See also Smolin [2017] who analyzes optimal evaluation policies for an agent who decides when to quit.

[2]Two other differences between our model and those two papers is that we have no discounting/delay costs and that we allow the sender to choose the order of actions, while in those papers, the order is determined by the passage of time.

one becomes a source of additional information about state two and vice versa. This is what we relate to as the information leakage problem. Unlike in Bergemann and Morris [2016b], our additional information is chosen optimally by the sender, and the sender may decide not to follow optimal persuasion for decision A in order to limit the indirect effect he creates for decision B.

Li and Norman [2018] consider sequential information provision by multiple senders. They focus on the strategic interactions among senders and analyze their implications for overall information provision. Thus, in Li and Norman [2018] the sender's ability to provide information is constrained by the presence of other senders, whose preferences are not necessarily aligned. In contrast, in our paper, the sender's ability to design information is constrained because the state of the world has a correlation structure.

Finally, our paper is also related to the literature on multi-product firms (Gamp [2016]) and products with multiple characteristics (Turlo [2018]), where the sender determines the order for information acquisition. In our paper, the sender decides the order in which the receiver takes her actions.

2 The Model

There are two players, a sender and a receiver. There are n binary states $\{\omega_i\}_{i=1}^n$, where $\omega_i \in \{0, 1\}$. The common prior belief regarding these states need not be symmetric or independent. We denote by x_i the prior probability of $\omega_i = 1$. When the receiver observes signal S (that can be correlated with all or some states), he updates his beliefs about the n states. Let $x_{s,i}$ be the updated belief about ω_i conditional on observing $S = s$.

The receiver takes n actions $\{a_i\}_{i=1}^n$. The receiver's action a_i depends on the belief that $\omega_i = 1$. We denote his i'th action by the function $a_i(x)$.[3]

[3] Since we are focusing on the sender's payoffs, we do not specify the receiver's payoffs for which these actions are optimal. How the optimal actions change with beliefs depends on the receiver's joint preferences over actions and states.

The sender's utility as a function of the action i is given by $f_i(a_i)$. The sender's utility as a function of the expected action profile, $\{a_i\}_{i=1}^n$, is additive in the different actions and is given by:

$$\sum_{i=1}^{n} f_i(a_i).$$

With some abuse of notation we sometimes refer to f_i also as a function of x_i; that is, $f_i(x_i) \equiv f_i(a_i(x_i))$. Finally, we let $cav f$ denote the concavification of function f.

We define two types of persuasion procedures: a simultaneous and sequential one.

Definition 1. *A simultaneous procedure is described by a single signal S that induces vector of posterior beliefs $x_i(s)$. For that procedure, the sender's expected payoff is:*

$$E_S[\sum_{i=1}^{n} f_i(x_{s,i})].$$

Definition 2. *A sequential procedure is described by a sequence of signals $\{S_i\}_{i=1}^n$ and a permutation $\pi : \{1..n\} \to \{1..n\}$ specifying the order of receiver decisions. In step i the sender sends signal S_i and the receiver takes action $a_{\pi(i)}$ (as a function of the posterior beliefs generated by all the signal realizations observed so far). Signal S_i may depend on the realizations of earlier signals, $\{s_j\}_{j=1}^{i-1}$.*

An important assumption about the sequential procedure is that once the agent takes action i, he cannot change it later even if new information comes to light and he regrets his earlier decision.

In the next section we compare these two procedures. Note that in the sequential procedure, the permutation (the order in which the receiver makes decisions) is chosen once ex-ante. Later we also consider a more flexible dynamic procedure, where the sender can adjust the order of decisions after each realization of the sequence of signals. However, we consider the procedure with a predetermined order of decisions to be more realistic in many applications.

Remark 1. *We interpret the game as having one receiver taking multiple actions. Equivalently, we could interpret it as having multiple receivers, each taking one action. In the latter case, we assume that all communication between the sender and the receivers is public. Even though each receiver cares only about 'their' dimension of the state, when the states are correlated, they all pay attention to the whole signal. See also in the next section the discussion of the upper bound on seller's payoff,* FB.

Remark 2. *Throughout the paper, the sender cannot condition future signals on the actions taken by the agent. We do not allow future information to be used as an incentive to induce the receiver to take sub-optimal actions today. This way our results identify a distinctively different benefit of dynamic persuasion (reduction of information leakage) than the intuition of using the information as a carrot in Ely and Szydlowski [2020] and Orlov et al. [2020]. When the game is interpreted as having a sequence of receivers, not allowing the sender to condition future signals on current actions is without loss of generality. If there is only one receiver, allowing the sender to commit to conditioning future signals on past actions would further relax the problem.*

3 Simultaneous versus Sequential Persuasion

We start by establishing an upper bound on the sender's payoffs:

Lemma 1. *For any procedure (simultaneous or sequential) the sender's payoff is at most*

$$FB(\{x_i\}) \equiv \sum_{i=1}^{n} cav f_i(x_i).$$

As we know from (Kamenica and Gentzkow 2011) $cav f_i(x_i)$ equals the optimal payoff that the sender could obtain from action i if it was the only action it had to persuade. Hence, FB, which stands for first best, represents the sum of payoffs if the sender could optimize each dimension separately. It is an upper bound for any procedure, sequential or simultaneous, since no procedure can achieve for dimension i a payoff greater than $cav f_i(x_i)$.

Note that this payoff would be obtained by the optimal simultaneous procedure if the prior beliefs about the states were independent. That is, if we changed the prior to one defined by the following two properties:

- The marginal distribution for each state equals that of the original prior (i.e., the x_i's are the same).
- The distributions of the all states are independent of one another.

With this modified prior there is no information leakage: when the sender persuades the receiver about state ω_i he can do it without changing the beliefs about the other states. In contrast, providing information about state ω_i under the original prior reveals information about other states as well. This information leakage constrains the sender's ability to persuade optimally in every dimension at once. Therefore, as we saw in our examples in the introduction, when the correlation of states is strong enough the sender may not be able to achieve the first-best payoff with either procedure.

Remark 3. *The first-best payoff could be also achieved in a model with n receivers and secret persuasion - each agent observing a private signal and taking decisions without observing the signal realizations or actions of others. Hence, a possible interpretation of our results is that they describe the cost of persuasion of multiple receivers and how sequential procedures can mitigate those costs.*

3.1 Strict Improvement from Sequential Procedure

We now ask when a sequential procedure can strictly improve upon the simultaneous one. Note that the outcome of any simultaneous procedure can be replicated by a sequential procedure with all signals except for the first one being uninformative. So, trivially the optimal sequential procedure weakly dominates the simultaneous one. As we have seen in the examples above, sometimes the improvement is strict.

Our main result is that what we observed in the examples is not a coincidence: if there exists a scope for improvement (i.e., the simultaneous procedure cannot achieve the first-best bound), then the optimal sequential procedure strictly improves upon the simultaneous one:

Theorem 1. *The optimal sequential procedure strictly dominates (in terms of the sender's payoffs) the optimal simultaneous procedure if and only if the latter fails to achieve the first-best payoff,* $FB(\{x_i\})$*.*

Proof. The 'only if' direction follows from the fact that $FB(\{x_i\})$ is an upper bound on both procedures. So we only need to prove the 'if' part. Consider a simultaneous procedure with signal S that fails to achieve the first-best outcome, $FB(\{x_i\})$. This implies that for at least one action, i:

$$E_S[f_i(x_{s,i})] < cav f_i(x_i).$$

We shall argue that there exists a sequential mechanism that improves the payoff from action i while keeping the payoffs from all other actions unchanged. Our proof is based on considering three cases:

- Case 1: $f_i(x_i) = cav f_i(x_i)$.
  In this case the expected payoff from action i is strictly lower when the signal S is revealed, compared to case where it does not. Define the sequential procedure as follows: in the first step the receiver takes action a_i based solely on the prior. In the second step, signal S (from the

simultaneous procedure) is revealed (and there is no more information revealed later). Then the receiver takes the remaining actions in an arbitrary order.

Following case 1, we can assume that $cav f_i(x_i) > f_i(x_i)$. Based on this, there exist $x_i^L, x_i^R \in [0,1]$ such that $x_i^L < x_i < x_i^R$ and payoffs satisfy $cav f_i(x_i^L) = f_i(x_i^L)$, $cav f_i(x_i^R) = f_i(x_i^R)$, $x_i = \lambda \cdot x_i^L + (1-\lambda) \cdot x_i^R$, and $cav f_i(x_i) = \lambda \cdot f_i(x_i^L) + (1-\lambda) \cdot f_i(x_i^R)$. Moreover, for any $x_i \in (x_i^L, x_i^R)$, $f_i(x_i) < cav f_i(x_i)$.[4]

- Case 2: For some realizations $S = s$, we have $x_{s,i} \in (x_i^L, x_i^R)$.
  Since $f_i(x_{s,i}) < cav f_i(x_{s,i})$ we construct a sequential mechanism where action i is the last step. All other actions are taken based on the original signal S. In the last step, conditional on observing $S = s$ the procedure provides more information regarding state i, which induces posteriors of x_i^L and x_i^R for state ω_i.

- Case 3: For all realizations $S = s$ we have $x_{s,i} \notin (x_i^L, x_i^R)$.
  Now, in the the sequential procedure, action i is taken first. This action is based on a signal S', which is a garbling of S. By Lemma 2, there exists a garbling S' of S which has two realizations: s'_1, s'_2 so that $x_{s'_1,i} = x_i^L$ and $x_{s'_2,i} = x_i^R$. We construct a sequential procedure where action i is taken first after S' is realized. The original S is revealed in the second step, and the receiver takes all other actions. While the original signal does not achieve $cav f_i(x_i)$, this new sequential procedure does (and leaves the payoffs for all other dimensions unchanged).

□

[4]Generically, x_i^L, x_i^R are unique. In case they are not, we define them to be the closest points to x_i.

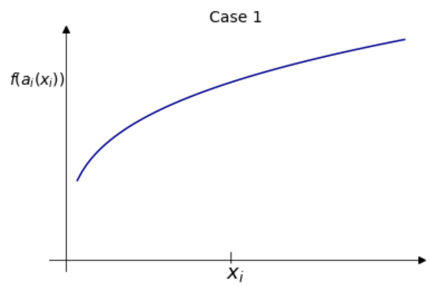

Figure 1: Proof of Theorem 1, Case 1

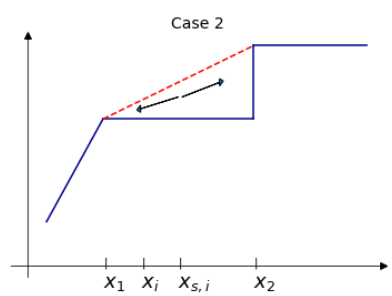

Figure 2: Proof of Theorem 1, Case 2

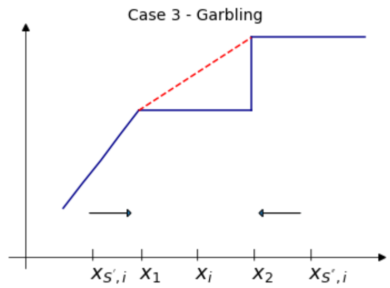

Figure 3: Proof of Theorem 1, Case 3

Lemma 2. *Suppose that X is a random variable and there are two values $x_1 < x_2 \in R$ such that $Pr(X < x_1) > 0, Pr(X > x_2) > 0, Pr(X \in (x_1, x_2)) = 0$ and $E[X] \in (x_1, x_2)$. Then there exists a binary random variable X' with realizations x_1', x_2' which is a garbling of X where $E[X|X' = x_1'] = x_1, E[X|X' = x_2'] = x_2$.*

3.2 Two Purchase Decisions

To understand better the conditions under which the simultaneous procedure achieves the first-best payoff, we consider a setup that is based on our examples in the introduction with two binary actions (buy, not buy). We assume that the threshold belief for buying is the same for both actions: α. We allow for imperfect correlation of the states. We assume without loss of generality that $x_1 \geq x_2$ and focus on the case where $x_2 \leq x_1 < \alpha$ (so the seller would like to persuade both actions).

To describe the prior belief, we introduce the following notation. The overall state of the world is $\omega = (\omega_1, \omega_2)$. It can take the following values: $\omega_{ij} = (\omega_1 = i, \omega_2 = j)$ and the prior belief is $x_{ij} = Pr(\omega = \omega_{ij})$. For example, x_{10} is the prior probability that $\omega_1 = 1$ and $\omega_2 = 0$.

In this setup our result is that the simultaneous procedure fails to achieve the first-best outcome, if and only if $\alpha > \frac{1}{2}$ and ω_1 and ω_2 are sufficiently strongly negatively correlated. Formally we argue that:

Proposition 1. *The simultaneous procedure achieves the first-best outcome if and only if one of the following two inequalities hold:*

(*i*) $x_{11} \geq x_{01}\frac{2\alpha-1}{1-\alpha}$, *or*

(*ii*) $x_{10} + x_{01} \leq x_{00}\frac{\alpha}{1-\alpha} + 2x_{11}\frac{1-\alpha}{2\alpha-1}$.

While we provide a formal proof in the Appendix, we shall next introduce the logic behind it. A simultaneous procedure is equivalent to choosing a joint distribution of signal S and state of the world ω. A simultaneous procedure achieves the first-best outcome if and only if the support of the marginal distribution of the posterior ω, conditional on $S = s_{ij}$, is $\{0, \alpha\}$.

Consider the following joint distribution of S and ω for some parameters $\{b, c, d, e, g\}$:

$S \backslash \omega$	ω_{11}	ω_{10}	ω_{01}	ω_{00}
s_{11}	x_{11}	b	b	c
s_{10}	0	$x_{10} - b$	0	d
s_{01}	0	0	$x_{01} - b$	e
s_{00}	0	0	0	g

If there exists a joint distribution that achieves the first-best payoff, then there exists a signal S that has distribution of this form (for appropriate values of of the parameters) that also achieves the first-best payoff. To see this, note that if it is possible to achieve the first-best payoff, it is sufficient to do it with a signal S that has at most four realizations that correspond to the posteriors that induce the buyer to take one of the four action combinations. He buys both products if the signal is s_{11}, only product A if the signal is s_{10}, only product B if the signal is s_{01}, nothing if the signal is s_{00}. For the first-best payoff it has to be that when he buys nothing his posterior belief is that the state is for sure ω_{00}, which explains the necessity of the last row of this joint probability distribution. Similarly, when he buys only one of the goods, he has to have a belief that that good is of high quality with probability α and the other good is of high quality with probability 0, which explains the other zeros in the table.

Finally, consider the joint probability $Pr(s_{11} \cap \omega_{11})$ (the top left entry in the table). The reason it must be equal to x_{11} is that when the state is ω_{11} in the first best outcome, the buyer buys both products for sure. If that entry were not equal to x_{11}, there would be a signal realization after which the buyer would not buy at least one of the goods, despite assigning a positive probability to both products being of high quality. Such a procedure would fail to achieve the first best.

With this observation, the rest of the proof of Lemma 1 consists of identifying conditions in terms of α and the prior distribution of ω for which there exist parameters $\{b, c, d, e, g\}$ such that such a joint distribution is feasible and achieves the first-best outcome.

4 Flexible Sequential Persuasion

So far, we have assumed that the sender predetermines the order of actions in a sequential procedure. This section considers sequential procedures where the order of actions is contingent on the realized signals. We name such procedures "flexible sequential" and show that the optimal procedure of that kind can always achieve the first-best payoff, FB.

To illustrate this, consider Example 1 from the introduction. The seller can achieve the first-best payoff by employing the following persuasion strategy. In the first step, the seller induces two posterior beliefs that A is of high quality: 0.8 or 0.2. Given the perfect negative correlation, the signal induces a posterior belief of 0.2 and 0.8 respectively, that B is of high quality. The buyer is then asked to buy the good with the higher posterior. Suppose the buyer buys the good A in the first step. After that, there is still a 20%

chance that good B has high quality. The seller then induces a posterior of 0.8 that B has high quality with a probability of 0.25. Since he sells with certainty in the first step and with probability 0.25 in the second step, his expected payoff equals 1.25, which is the first-best payoff.

A flexible sequential procedure is a sequential procedure where the order actions can depend on the value of realized signals. Formally, we define it as:

Definition 3. *A flexible sequential procedure is a sequence of k signals $\{S_i\}_{i=1}^k$ and a conditional action order $\pi(S_1...S_i) \in \{0, 1, ...n\}$, that specifies sequence of receiver actions conditional on the signals realized so far. In step i, conditional on the realizations of signals $S_1, ...S_{i-1}$, the sender sends signal S_i and the receiver takes action $a_{\pi(S_1...S_i)}$. When $\pi(S_1...S_i) = 0$ the receiver takes no action in step i. The number of steps, k, can be random but is finite almost surely.*

Implicit in this definition is that for a flexible sequential procedure to be well-defined, each dimension must eventually be chosen. That is, for every possible sequence of signal realizations, the conditional order $\pi(S_1...S_k)$ has to select each dimension once. Also, note that we allow for $k \geq n$. That is, we allow the sender to choose a procedure where the receiver takes no action in some steps.

4.1 The Pacman procedure

An optimal flexible sequential procedure can be described as a 'Pacman procedure.' We shall assume that for all i we have that $cav f_i(x_i) > f_i(x_i)$. If this is not the case, and $cav f_i(x_i) = f_i(x_i)$, then the optimal sequential procedure starts with dimension i and the receiver is asked to take action i without receiving any information.

Recall from the proof of Theorem 1 that x_i^L, x_i^R denote the optimal concavification beliefs for action i (they depend on the prior, but we suppress that notation). Let $x_i^L < x_i < x_i^R$. The Pacman procedure that achieves the first-best payoff is based on a sequence of signals and a contingent sequence of actions with the following property. The receiver takes action i only when $x_{s,i} \in \{x_i^L, x_i^R\}$ (we abuse notation by writing $x_{s,i}$ as the posterior belief given the realized sequence of signals). The idea is to construct a dynamic Bayesian persuasion procedure where, as we disclose information about some dimensions, the belief regarding all states $i \in \{1, ..., n\}$ remains in every step in the interval $[x_i^L, x_i^R]$. When the belief hits one of the boundaries of these intervals, the receiver is asked to take action a_i (and we remove that dimension from the continuation problem). This ensures that the procedure achieves the first-best payoff, provided that it ends in a finite time because from the ex-ante perspective, beliefs are a martingale, and actions are taken exactly at the concavification points.

For any i, consider some signal S that is correlated with ω_i and conditionally on ω_i independent of any other ω_j. Suppose that the signal S is rich in the sense that for every $p \in [0, 1]$ there exists s such that $x_{s,i} = p$. Finally for that signal define $h_{i,j}(p) \equiv Pr(\omega_j = 1 | x_{s,i} = p)$. This function describes how sensitive the beliefs about dimension j are in response to signals about dimension i. The following technical lemma shows that this sensitivity is at most one:

Lemma 3. *$|h'_{i,j}(p)| \leq 1$.*

To construct an optimal flexible sequential procedure (that achieves the first-best payoff FB), define:

$$\Delta_i = \min\{x_i^H - x_i, x_i - x_i^L\}$$
$$\Delta = \min_i \Delta_i$$
$$i^* = \arg\min_i \Delta_i.$$

Without loss of generality, assume that the minimization in Δ_{i^*} is achieved by $x_{i^*}^R$ (that is, $x_{i^*}^R - x_{i^*} \leq x_{i^*} - x_{i^*}^L$).

We now construct the Pacman procedure that approximates in finitely many steps the FB. In step 1 of the procedure, we choose a signal conditional on ω_{i^*} (and conditionally independent of the other states) that generates posterior beliefs about ω_{i^*} of $x_{i^*} \pm \Delta$, each with probability 0.5. Consider the posteriors for all other states and denote the posterior for ω_i by $x_{s,1}$. Lemma 3 implies that for both realizations of the signal in step 1:

$$\forall i, s : x_{s,i} \in [x_i^L, x_i^R].$$

Conditional on the posterior for ω_{i^*} being $x_{i^*} + \Delta = x_{i^*}^R$, we ask the receiver to take action i^* (otherwise we tell the receiver not to take any actions). We then iterate on this procedure. If action i^* is taken, we remove it from the consideration set. Given the posteriors from the previous step, we redefine Δ and i^* and repeat.

To see that this procedure achieves the first-best payoff, note that in each step, the receiver acts in one of the remaining dimensions with probability 0.5 and takes actions only at the optimal concavification thresholds. It implies that the procedure ends in finite time almost surely. This proves the following theorem:

Theorem 2. *The Pacman procedure achieves the first-best payoff.*

While we construct the Pacman procedure to take more rounds than the number of dimensions, it is possible to reduce the number of steps. Namely, we can create a more complex sequence of signals to take n steps and achieve the first-best payoff exactly. The intuition is that when the first signal realization moves the beliefs to $x_{i^*}^L$, the sender does not need to reveal the signal. Instead, he can automatically compute the next optimal signal and reveal it jointly with the first signal realization. If the realization of the second signal also does not move beliefs to any of the concavification thresholds, we iterate. In the limit, we obtain a random variable such that every realization moves the beliefs to at least one of the optimal concavification thresholds without moving any beliefs outside the $[x_i^L, x_i^R]$ region. So in every step, one of the actions is taken, and information leakage never moves beliefs 'too far.'

Finally, the Pacman procedure helps us understand an optimal sequence of persuasion. With this procedure, it is optimal to rank decisions based on the distance between the prior and the nearest concavification threshold (the Δ_i's) and then persuade based on this order. The only caveat is that the ranking of Δ_i's can change due to the disclosed information, so the order of persuasion depends not only on the ex-ante parameters but also on the information learned in the process of persuasion. The other takeaway from this construction is that with a flexible procedure, the order of information matters less than disclosing information gradually and then choosing the proper order of actions. The most important thing about this procedure is that information is disclosed gradually to ensure that a concavification threshold is not jumped over in any step.

5 Simple versus Flexible Sequential Persuasion and Optimal Order of Persuasion.

We now turn to when the optimal flexible sequential procedure can strictly improve upon the simple sequential procedure we introduced in the two previous sections. Since the flexible sequential procedure always achieves the first-best payoff, the question is under what conditions the simple sequential procedure can achieve it too. In order to provide a partial answer to this question, we start with the special case of perfectly correlated states. That was the case in Examples 1 and 2 in the Introduction and, in general, makes it the hardest to achieve the first-best payoff. Analyzing this case also sheds some light on the optimal order of persuasion.

5.1 Perfectly Correlated States

Let x denote the (common) prior belief that all states are $\omega_i = 1$. Without loss of generality, we assume perfect *positive* correlation. If some states are perfectly negatively correlated with state 1, then we switch the definition of those states, inducing a positive correlation. For example, in Example 1, we flip the meaning of the state for good B to mean that state 1 represents low quality. After that transformation, the buyer purchases good B if and only if the belief about state B is below 0.2.

The payoff of the optimal simultaneous procedure is:

$$U^*_{Sim}(x) = cav(\sum_{i=1}^{n} f_i)(x). \tag{1}$$

The payoff from a sequential procedure when there are two actions and the receiver takes first action 1 and then action 2 is $cav(f_1 + cav f_2)(x)$. The intuition is that if the sender induces in the first stage posterior $\hat{x}$, in the second stage, the persuasion payoff is $cav f_2(\hat{x})$. Taking this into account, the sender realizes that the total payoff as a function of the first-period posterior is $(f_1 + cav f_2)(\hat{x})$. This, in turn, is maximized in stage 1 by the standard concavification at the prior, yielding $cav(f_1 + cav f_2)(x)$.

For more states and some permutation $\pi : \{1..n\} \to \{1..n\}$, the optimal sequential procedure that follows this order of actions yields a payoff that can be defined recursively:

$$H_n(\pi) \equiv cav f_{\pi(n)}$$
$$H_i(\pi) \equiv cav(f_{\pi(i)} + H_{i+1}(\pi)).$$

The payoff of the optimal sequential procedure given prior x can therefore be expressed as:

$$U^*_{Seq}(x) = max_\pi H_1(\pi)(x). \tag{2}$$

The result in Theorem 1 in the case of perfect correlation of states can be viewed as the comparison of (1) to (2)). In particular, whenever $U^*_{Sim}(x) < FB(x)$ then $U^*_{Sim}(x) < U^*_{Seq}(x)$.

This characterization can be used to show that the sequential procedure in Example 1 is optimal despite not achieving the first-best payoff (see the Appendix).

Our goal in this section is to examine when the simple sequential procedure achieves the first-best payoff. Define

$$\Delta_i^L \equiv x - x_i^L$$
$$\Delta_i^H \equiv x_i^H - x,$$

where recall that x_i^L and x_i^H are the optimal concavification thresholds to the left and to the right of the prior for dimension i (given the prior). Let π^L be the list of $i's$ in the order of Δ_i^L and analogously define π^H. For example if $n = 3$ and the vectors are $\Delta^L = (0.1, 0.3, 0.2)$ and $\Delta^H = (0.4, 0.1, 0.2)$, then $\pi^L = (1, 3, 2)$ and $\pi^H = (2, 3, 1)$. In words, π^L and π^H are the orderings of the dimensions in terms of the ranking of the optimal concavification thresholds to the left and to the right of the prior, respectively. Note that if there are ties in some Δ_i's (see Example 2), these orderings are not unique (all ways of resolving these ties are allowed).

Proposition 2. *If states are perfectly (positively) correlated, the optimal sequential procedure achieves the first-best payoff if and only if there exist orderings such that $\pi^L = \pi^H$ (i.e., for at least one resolution of ties).*

When this condition is satisfied, the optimal sequence of persuasion follows these orderings: the sender persuades first about the good with the smallest Δ_i^H, then the second smallest, and so on.

Proof. Suppose this condition is satisfied. That means that the concavification thresholds are nested around the prior. Moving in the proposed sequence, in every step, we reach one of the two concavification thresholds for one of the dimensions, eliminating them deterministically one by one. Specifically, let $i = \pi^L(1)$. In step 1 choose a binary signal S_1 inducing posteriors $x_s \in \{x_i^L, x_i^H\}$. Let the receiver take action i and remove that dimension from the consideration set. Next, redefine $i = \pi^L(2)$ and again choose a binary signal that induces posteriors $x_s \in \{x_i^L, x_i^H\}$. Since the thresholds are nested, in step $k+1$ the posteriors remain interior to $[x_j^L, x_j^H]$ for every dimension j later in the sequence than dimension $\pi^L(k)$. That allows the sender to induce posterior beliefs equal to the corresponding optimal thresholds in every step, proving that this order achieves the first-best payoff.

If the two orderings are not the same, then this implies that there are two dimensions i, j such that $\Delta_i^L < \Delta_j^L$ but $\Delta_i^H > \Delta_j^H$ (as in Example 1). In that case it is not possible to achieve the first-best payoff with the simple sequential procedure, since if the sender chooses an order of actions such that the receiver has to take action i before action j (without loss of generality) then to achieve the first-best payoff for dimension i we need to induce with positive probability the belief $x_i^H > x_j^H$. That in turn means that we do not obtain $cav f_j(x)$ on dimension j. $\square$

5.2 Imperfectly Correlated States

Proposition 2 allows us to also provide intuition about the imperfectly correlated case. Given priors x_i define generalize the definition:

$$\Delta_i^L \equiv x_i - x_i^L$$
$$\Delta_i^H \equiv x_i^H - x_i,$$

and define the orderings $\pi^L = \pi^H$ as before.

First, suppose that all the states are positively correlated (in the sense that all $h'_{i,j}(p)$ described in Lemma 3 are non-negative). If the orderings of dimensions in terms of ranking of Δ_i^L and Δ_i^H are the same ($\pi^L = \pi^H$), then the simultaneous procedure achieves the first-best payoff for all priors and (positive) correlations. Furthermore, if the orderings are not the same, no sequential procedure can reach the first-best payoff for a sufficiently high correlation. The intuition is the same as in the proof of Proposition 2. We can start with the dimension with the smallest Δ_i and induce the concavification thresholds. Since the $\Delta_i's$ are nested, we do not "jump over" the concavification thresholds for other dimensions even with perfect correlation. When correlation is imperfect, the posteriors on the other dimensions are even closer to the starting priors.

Second, when the correlations are negative, having the dimensions ranked in the same order in both directions is not enough.[5] The reason is that when we move x_i to x_i^H, even though we are guaranteed not to overshoot any x_j^H, we may overshoot some x_j^L. In that case, a sufficient condition for the sequential procedure to achieve the first-best payoff for all correlations is that for all i, j, if $\Delta_i^H < \Delta_j^H$, then both $\Delta_i^L < \Delta_j^L$ and $\Delta_i^H < \Delta_j^L$. That condition is also tight in the following sense. If one of the inequalities is reversed, for some correlation structure (positive or negative), the optimal sequential procedure does not achieve the first-best payoff.

6 Conclusions

We have characterized the benefits of sequential persuasion of a receiver who takes multiple irreversible actions with arbitrary correlations between binary states and arbitrary sender's payoffs. Correlations between states constrain the sender in the sense that when he sends information about one of the states, he necessarily reveals some information about other states. This information leakage may prevent the sender from receiving the first-best payoff that he could obtain if the states were independent.

Our first main result is that whenever the information leakage problem

[5]Note that if there are more than two dimensions and general imperfect negative correlation, it may be impossible to change the definition of states to assure that all pairwise correlations are positive.

is harmful to the sender, the sender can do strictly better with some simple sequential procedure. This benefit of dynamic persuasion is distinct from the information-as-a-carrot benefit stressed in the earlier literature. Our second main result is that if the seller could use a flexible sequential procedure, he can always completely mitigate the information leakage problem. Finally, we provide intuition about the optimal sequencing of the persuasion. We provide a mathematical foundation for the heuristic that optimal persuasion starts with the action that requires the least persuasion (in the sense of moving the beliefs by the smallest amount).

Admittedly, in some situations opposing forces not captured by our model could push in favor of one-time, simultaneous persuasion. One such force is that taking decisions could reveal additional information to the receiver. For example, if in Example 2, the receiver learns the value of the good immediately upon purchase, then sequential persuasion can make the information leakage problem worse.

7 Appendix

Proof that the sequential procedure in Example 1 is optimal.

Proof. As we discussed in Section 5.1, when the states are perfectly correlated, the highest payoff from sequential persuasion is

$$U^*_{Seq}(x) = max\{cav(f_1 + cav f_2)(x), cav(f_2 + cav f_1)(x)\},$$

Since in this problem the two actions are symmetric, we need only consider a single order. We have

$$\begin{aligned}(f_1 + cav f_2)(x) &= 1 + \frac{x}{0.8} \text{ if } x \leq 0.2 \\ &= \min\{\frac{x}{0.8}, 1\} \text{ if } x > 0.2\end{aligned}$$

Concavification of that function gives us:

$$\begin{aligned} cav(f_1 + cav f_2)(x) &= 1 + \frac{x}{0.8} \text{ if } x \leq 0.2 \\ &= 1 + \frac{1}{4}\frac{1-x}{0.8} \text{ if } x > 0.2 \end{aligned}$$

In particular, if we start with a prior $x = \frac{1}{2}$ the highest sequential payoff is

$$U^*_{Seq}(p) = 1 + \frac{1}{4}\frac{1-0.5}{0.8} = 1.1563.$$

Proof of Lemma 2

We first argue that without loss of generality we can consider binary signals S with only two realizations a and b where $a \leq x_1 < x_2 \leq b$. This follows from the fact that we can define a signal $\hat{S}$ where $\hat{S} = E(S|S \leq x_1)$ when $S \leq x_1$ and $\hat{S} = E(S|S \geq x_2)$ where $S \geq x_2$. Since, $\hat{S}$ is garbling of S, garbling of $\hat{S}$ is also garbling of S.

Consider a signal S_z indexed by z; the signal is binary with realizations $\{s_1, s_2\}$ that occur with probabilities $\{1-z, z\}$, respectively. Specifically, s_2 occurs with probability $\alpha \cdot z$ when $S = a$ and with probability $(1-\alpha) \cdot z$ when $S = b$ where $\alpha \in [0,1]$ is defined by:

$$x_2 = (1-\alpha) \cdot a + \alpha \cdot b$$

As a result we have that $\forall z : E(S|S_z = s_2) = x_2$. The feasible range for z is $[0, min\{\frac{Pr(S=x_2)}{\alpha}, \frac{Pr(S=x_1)}{1-\alpha}\}]$. We first note that $\frac{Pr(S=x_2)}{\alpha} \leq \frac{Pr(S=x_1)}{1-\alpha}$. This implies that the feasible range for z is $[0, \frac{Pr(S=x_2)}{\alpha}]$. To see why this holds note that:

$$\frac{Pr(S = x_2)}{\alpha} > \frac{Pr(S = x_1)}{1-\alpha} \Rightarrow \frac{Pr(S = x_2)}{Pr(S = x_1)} > \frac{\alpha}{1-\alpha}$$

This would imply that $E(S) > x_2$ which is a contradiction. Consider the other realization s_1 and let $H(z) \equiv E(S|S_z = s_1)$. The claim follows from the intermediate value theorem as $H(0) = E(S)$ and $H(\frac{Pr(S=x_2)}{\alpha}) = a < x_1$. Hence, we conclude that there exists z^* so that $H(z^*) = E(S|S_z^* = s_1) = x_1$. $\square$

Proof of Proposition 1

Proof. 'If' direction.

Say condition (i) $x_{11} \geq x_{01}\frac{2\alpha-1}{1-\alpha}$ holds. Then the following signal is feasible: $b = x_{01}$, $c = x_{11}\frac{1-\alpha}{\alpha}+x_{01}\frac{1-2\alpha}{\alpha}$, $d = [x_{10}-p_{01}]\frac{1-\alpha}{\alpha}$, $e = 0$ and $g = x_{00}-c-d$.

First we show that the suggested signal satisfies the obedience constraints.

$$x_{s_{11},i} = \frac{x_{11} + x_{01}}{x_{11} + x_{01} + x_{01} + x_{11}\frac{1-\alpha}{\alpha} + x_{01}\frac{1-2\alpha}{\alpha}} = \frac{x_{11} + x_{01}}{(x_{11} + x_{01})\frac{1}{\alpha}} = \alpha \tag{3}$$

$$x_{s_{10},1} = \frac{x_{10} - x_{01}}{(x_{10} - x_{01}) + (x_{10} - x_{01})\frac{1-\alpha}{\alpha}} = \frac{x_{10} - x_{01}}{(x_{10} - x_{01})\frac{1}{\alpha}} = \alpha \tag{4}$$

Equation (3), for example, shows that given the signal realization s_{11} expectation for dimension i is α, for $i = 1, 2$.

Now we show that the suggested signal is also a feasible joint distribution. To do this we must show that $c + d + e \leq x_{00}$. Expressing x_{00} as a complementary probability, yields $x_{00} = 1 - x_{11} - x_{10} - x_{01}$.

Therefore, $c + d + e \leq x_{00} \iff$

$$x_{11}\frac{1-\alpha}{\alpha} + x_{01}\frac{1-2\alpha}{\alpha} + [x_{10} - x_{01}]\frac{1-\alpha}{\alpha} \leq 1 - x_{11} - x_{10} - x_{01}. \tag{5}$$

This expression simplifies to:

$$x_{11}\frac{1}{\alpha} + x_{10}\frac{1}{\alpha} \leq 1, \tag{6}$$

which holds, since we assume that $x_i \leq \alpha$.

Now suppose that (i) is violated, but condition $(ii) : x_{10}+x_{01} \leq x_{00}\frac{\alpha}{1-\alpha}+2x_{11}\frac{1-\alpha}{2\alpha-1}$ is satisfied. Then the following signal is feasible: $b = x_{11}\frac{1-\alpha}{2\alpha-1}$, $c = 0$, $d = (x_{10} - x_{11}\frac{1-\alpha}{2\alpha-1})(\frac{1-\alpha}{\alpha})$ and $e = (x_{01} - x_{11}\frac{1-\alpha}{2\alpha-1})(\frac{1-\alpha}{\alpha})$.

First, we show that obedience constraints are satisfied.

$$x_{s_{11},i} = \frac{x_{11} + x_{11}\frac{1-\alpha}{2\alpha-1}}{x_{11} + x_{11}\frac{1-\alpha}{2\alpha-1} + x_{11}\frac{1-\alpha}{2\alpha-1}} = \alpha \tag{7}$$

$$x_{s_{10},1} = \frac{x_{10} - x_{11}\frac{1-\alpha}{2\alpha-1}}{x_{10} - x_{11}\frac{1-\alpha}{2\alpha-1} + (x_{10} - x_{11}\frac{1-\alpha}{2\alpha-1})\frac{1-\alpha}{\alpha}} = \frac{1}{1+\frac{1-\alpha}{\alpha}} = \alpha \tag{8}$$

$$x_{s_{01},2} = \frac{x_{01} - x_{11}\frac{1-\alpha}{2\alpha-1}}{x_{01} - x_{11}\frac{1-\alpha}{2\alpha-1} + (x_{01} - x_{11}\frac{1-\alpha}{2\alpha-1})\frac{1-\alpha}{\alpha}} = \frac{1}{1+\frac{1-\alpha}{\alpha}} = \alpha \tag{9}$$

Feasibility means that $d + e \leq x_{00}$, i.e.:

$$(x_{10} - x_{11}\frac{1-\alpha}{2\alpha-1})(\frac{1-\alpha}{\alpha}) + (x_{01} - x_{11}\frac{1-\alpha}{2\alpha-1})(\frac{1-\alpha}{\alpha}) \leq x_{00} \tag{10}$$

which is equivalent to condition (ii):

$$x_{10} + x_{01} \leq x_{00}\frac{\alpha}{1-\alpha} + 2x_{11}\frac{1-\alpha}{2\alpha-1}. \tag{11}$$

'Only if' direction.

If both conditions (i) and (ii) do not hold, then, for any feasible signal (S), at least one of the obedience constraints does not bind. This is so, because if all constraints bind, then $c+d+e > x_{00}$, i.e., the signal (S) is not a feasible joint distribution. We now prove this claim.

Since condition (i) is violated, we consider the following recommendation rule:

$b = x_{11}\frac{1-\alpha}{2\alpha-1}$, $c = 0$, $d = (x_{10}-x_{11}\frac{1-\alpha}{2\alpha-1})(\frac{1-\alpha}{\alpha})$ and $e = (x_{01}-x_{11}\frac{1-\alpha}{2\alpha-1})(\frac{1-\alpha}{\alpha})$.

We have shown in equations $7-9$, that the suggested signal satisfies the obedience constraints. The feasibility constraint becomes: $c + d + e \leq x_{00}$, which upon substituting for $c + d + e$ yields

$$(x_{10} - x_{11}\frac{1-\alpha}{2\alpha-1})(\frac{1-\alpha}{\alpha}) + (x_{01} - x_{11}\frac{1-\alpha}{2\alpha-1})(\frac{1-\alpha}{\alpha}) \leq x_{00}. \tag{12}$$

(12) is equivalent to condition (ii), which is a contradiction.

Now we show that there does not exist any other feasible signal that achieves the first-best payoff. Thus, we consider a signal with $c = c' > 0$. Then following is true:

$$c' + d' + e' > c + d + e > x_{00}, \tag{13}$$

where $c' > 0$ and d' and e' are the values of the suggested new signal and $c = 0$, $d = (x_{10} - x_{11}\frac{1-\alpha}{2\alpha-1})(\frac{1-\alpha}{\alpha})$ and $e = (x_{01} - x_{11}\frac{1-\alpha}{2\alpha-1})(\frac{1-\alpha}{\alpha})$. Condition (13) holds, because for any signal that achieves the first-best payoff, b is decreasing in c, whereas d and e are decreasing in b. Say $c = c' > 0$, then $b' < b = x_{11}\frac{1-\alpha}{2\alpha-1}$. If not, then

$$x_{s_{11},i} = \frac{x_{11} + x_{11}\frac{1-\alpha}{2\alpha-1}}{x_{11} + x_{11}\frac{1-\alpha}{2\alpha-1} + x_{11}\frac{1-\alpha}{2\alpha-1} + c'} < \alpha \tag{14}$$

and the obedience constraint for s_{11} is violated. If a recommendation rule achieves the first-best payoff, then $d' = (x_{10} - b')(\frac{1-\alpha}{\alpha})$ and $e' = (x_{01} - b')(\frac{1-\alpha}{\alpha})$, i.e., the obedience constraints for s_{10} and s_{01} bind. This completes the argument.

This means that if both conditions are violated, then there does not exist a feasible signal that induces a distribution of marginal posteriors with support $\{0, \alpha\}$. QED

□

Proof of Lemma 3

Proof. Note that:

$$\begin{aligned} h_{i,j}(p) &= p \cdot Pr(\omega_j = 1|\omega_i = 1) + (1-p) \cdot Pr(\omega_j = 1|\omega_i = 0) \\ &= Pr(\omega_j = 1|\omega_i = 0) + p \cdot [Pr(\omega_j = 1|\omega_i = 1) - Pr(\omega_j = 1|\omega_i = 0)] \end{aligned}$$

The proof then follows from $Pr(\omega_j = 1|\omega_i = 1), Pr(\omega_j = 1) \in [0, 1]$. □

References

Robert J Aumann, Michael Maschler, and Richard E Stearns. *Repeated games with incomplete information*. MIT press, 1995.

Dirk Bergemann and Stephen Morris. Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics*, 11 (2):487–522, 2016a.

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

Jeffrey C Ely. Beeps. *American Economic Review*, 107(1):31–53, 2017.

Jeffrey C Ely and Martin Szydlowski. Moving the goalposts. *Journal of Political Economy*, 128(2):468–506, 2020.

Tobias Gamp. Guided search, 2016.

Johannes Hörner and Andrzej Skrzypacz. Selling information. *Journal of Political Economy*, 124(6):1515–1562, 2016.

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

Fei Li and Peter Norman. Sequential persuasion. *Available at SSRN 2952606*, 2018.

Dmitry Orlov, Andrzej Skrzypacz, and Pavel Zryumov. Persuading the principal to wait. *Journal of Political Economy*, 128(7):2542–2578, 2020.

Alex Smolin. Dynamic evaluation design. *Available at SSRN 3051703*, 2017.

Sergey Turlo. Persuasion through ordered information. 2018.