Persuasion through Trial Design

Alesandro Arcuri

March 2021

Abstract

A researcher wants to persuade a policymaker to adopt his treatment, which may be either good or bad. The policymaker wants to adopt the treatment when it is good and not when it is bad. The researcher chooses how many subjects to enroll in a trial, under either the *sequential sampling* regime or the *pre-registration* regime. Each subject improves with probability ρ when the treatment is good and probability $1 - \rho$ when the treatment is bad. Under pre-registration, the researcher commits to his choice of sample size at the start of the trial, while under sequential sampling the researcher can observe each subject outcome before deciding whether to continue the trial or not. I show that under sequential sampling, as $\rho \rightarrow_{+} .5$, the researcher can achieve his first-best Bayesian persuasion outcome, which minimizes the policymaker's utility over attainable BP outcomes. I show that under pre-registration, however, when the good state is at at least as likely as the bad state ex ante, full revelation is optimal for the researcher, which maximizes policymaker utility over the set of attainable BP outcomes.

1 Introduction

Consider a researcher testing the efficacy of a new treatment (e.g. an educational intervention or a new pharmaceutical), and suppose the situation is as follows. The new treatment is either good or bad. When administered to a subject randomly drawn from some population of interest, if the treatment is good the subject will improve with probability $\rho \in (.5, 1)$, and if it is bad the subject will improve with probability $1 - \rho$.

The only control the researcher has over the test is the number of subjects to enroll. In the *pre-registration* regime, the sender commits to a sample size before any results are realized. In the *sequential sampling* regime, the sender observes the results of each subject before deciding whether to enroll the next one, and the trial ends when the sender declines to enroll any further subjects.

Under either regime, after the trial is over, the complete history (i.e. the status of each subject) is viewed by a policymaker. Depending on the outcome of the trial, the policymaker will make a binary adopt/reject decision regarding the treatment. The policymaker wants to adopt the sender's treatment when it is good, and reject it when it is bad. The policymaker is an expected utility maximizer, and her preferences can be parameterized by a cutoff belief z, such that she prefers to adopt the sender's product if and only if her belief that the state is good is at least z. The sender seeks to maximize the probability of policymaker adoption.

I show that under sequential sampling, as $\rho \to_+ .5$ the researcher's payoff approaches his first-best Bayesian persuasion payoff. At the same time the policymaker's payoff approaches her first-worst Bayesian persuasion payoff.

Under pre-registration, when $\mu \geqslant .5$, regardless of the values of z or ρ, I show that full revelation is optimal for the researcher. This outcome uniquely maximizes the policymaker's utility over all possible information structures.

Taken together, these results imply that when $\mu \geqslant .5$, as $\rho \to_+ .5$, the receiver's utility under sequential-sampling approaches her worst possible payoff under Bayesian persuasion, and under pre-registration approaches her best possible payoff under Bayesian persuasion. This immediately highlight the importance of pre-registration. The first result, (along with a similar continuous-time result from Morris and Strack (2019)) suggests that the Bayesian persuasion framework may be a good model for trial design under sequential sampling. The second result suggests that it may be much less useful for making predictions about trial design under pre-registration. Notably, the second result follows from the symmetry of subject outcomes, and would similarly hold even if subject outcomes were instead normally distributed.

Registration of clinical trials is already the norm in many instances. The International Committee of Medical Journal Editors summarize their policy as follows: "Briefly, the ICMJE requires, and recommends that all medical journal editors require, registration of clinical trials in a public trials registry at or before the time of first patient enrollment as a condition of consideration for publication. Editors requesting inclusion of their journal on the ICMJE website list of publications that follow ICMJE guidance should recognize that the listing implies enforcement by the journal of ICMJE's trial registration policy."[1] In contrast, allegations of publication bias are more common in fields like psychology without such commitments to pre-registration.

A 2015 study published in PLoS ONE looks at a sample of large clinical trials funded by the National Heart, Lung, and Blood Institute (NHLBI) from 1970 to 2012. Due to Federal legislation mandating study registration, "after the year 2000, all (100%) of large NHLBI [trials] were registered prospectively in ClinicalTrials.Gov prior to publication. Prior to 2000 none of the trials (0%) were prospectively registered." (Kaplan and Irvin 2015). The authors go on to find that " For trials published before the year 2000, we found that 17 out of 30 (57%) reported significant benefit for their primary outcome. In the new era where primary outcomes are prospectively declared (published post 2000), only 2 of 25 trials (8%) reported a significant benefit ($\chi^2 = 12.2$, p = 0.0005)" (Kaplan and Irvin 2015).

[1] http://www.icmje.org/recommendations/browse/publishing-and-editorial-issues/clinical-trial-registration.html

In this case, "registration" is done through the clinicaltrials.gov website, which is also one of the forms of registration the ICMJE recognizes. This requires researchers to register a target number of participants to be enrolled, as well as register their sampling method, ideally before enrolling their first subject.[2] However, it also requires researchers to register much more information, including the outcome measures of interest, and so it is hard to know how much of its impact is due specifically to making researchers register their sample sizes. Nevertheless, the model I present here suggests that forcing researchers to commit to their sample sizes may have a large impact on trial quality by itself.

The rest of the paper is structured as follows. The remainder of this section discusses related literature. Section 2 presents the model. Section 3 defines the sender's choice sets in the regimes of interest. Section 4 presents the main results. Section 5 concludes.

1.1 Related Literature

This paper is part of the recent literature on Bayesian persuasion and information design, following Kamenica and Gentzkow (2011). Brocas and Carillo (2007) consider a problem very similar to my model under the sequential sampling regime. Their sender similarly chooses when to stop generating symmetric Bernoulli signals in a binary-state environment. However, their receiver has access to three actions instead of two, they are not interested in the sender's welfare as $\rho \to_+ .5$, and they do not consider any kind of pre-registration. Morris and Strack (2019) also looks at a form of persuasion with sequential sampling. In their model, the sender chooses when to stop the flow of information in continuous time with a Brownian motion noise term. They show that by choosing an appropriate stopping rule, the sender can attain his first-best Bayesian persuasion utility, just as in my model as $\rho \to_+ .5$. However, they do not consider the case of pre-registration. Henry and Ottaviani (2019) also consider the problem of a sender who can choose when to stop the flow of information in continuous time, though they also do not consider pre-registration either.

2 Model

The state of the world ω is a random variable distributed over $\Omega = \{0, 1\}$, according to a commonly known interior prior $\mu = Pr(\omega = 1) \in (0, 1)$. There is a receiver ("she") who has access to an action set $A = \{0, 1\}$, where action 1 can be thought of as adopting some treatment and 0 as rejecting it. Before she makes a decision, a sender ("he") chooses an information structure π, which stochastically maps the state of the world into some outcome space, from some exogenously given set Π. After the receiver views the sender's choice of π and its outcome, she chooses her action a, and payoffs $v(a, \omega)$ and $u(a, \omega)$ of the sender and receiver respectively are realized. I assume that the receiver (i) strictly prefers to match her action to the state, and (ii) strictly prefers action 0 at

[2] https://prsinfo.clinicaltrials.gov/definitions.html

the prior μ, and I normalize her utility to be $u(0,\omega) = 0, u(1,\omega) = \omega - z$ for $\omega \in \{0,1\}$, for some $z \in (\mu, 1)$. I assume the sender's utility is determined only by whether the audience adopts the sender's product, $v = a$.

The key difference between this model and a standard Bayesian persuasion problem is the restrictions on the sender's choice set in either regime. I give a formal definition of Π for the two regimes of interest in the next section.

3 Information Structures

A signal $\pi = (S, \pi_0, \pi_1)$ consists of an outcome space S, and a pair of probability distributions (π_0, π_1) over S. The outcome s of π is distributed according to π_ω over S when the state is ω. I will refer to the universal set of all possible such information structures as Π^*.

It will also be useful to think of signals in terms of their induced conditional distributions over the receiver's actions. Let $p(\pi) = (p_\omega(\pi))_\omega = (Pr_\pi(a = \omega|\omega))_\omega$, so that $p_\omega(\pi)$ is the probability that the receiver chooses her preferred action in state ω given π. As an example, the conditional action distribution induced by full revelation is $(p_0(\pi^{full}), p_1(\pi^{full})) = (1,1)$, and the distribution induced by no revelation is $(p_0(\pi^{no}), p_1(\pi^{no})) = (1,0)$, since $z > \mu$. Define $P(\Pi) = \{p(\pi) : \pi \in \Pi\}$, so that $P(\Pi) \subset [0,1]^2$ is the set of action distributions that can be induced by a sender with choice set Π. With a slight abuse of notation, I will also write the expectations of receiver and sender payoffs given $p(\pi)$ as $u(p) = \mu(1-z)p_1 - (1-\mu)z(1-p_0)$ and $v(p) = \mu p_1 - (1-\mu)p_0$ respectively. Then instead of having the sender choose a trial design from Π, we may equivalently consider his choice to be over action distributions from $P(\Pi)$.

3.1 The Unrestricted Choice Set $P(\Pi^*)$

Suppose the sender can choose any signal $\pi \in \Pi^*$. Then his problem is as described in Kamenica and Gentzkow (2011). He can restrict his attention to signals which have two outcomes, one of which, s_A, induces the receiver to adopt, and one of which, s_R, induces the receiver to reject. The receiver will find such a signal incentive-compatible if the expected value of adopting after seeing outcome s_A is non-negative. Mathematically, this requires that $\frac{\mu\pi_1(s_A)}{\mu\pi_1(s_A)+(1-\mu)\pi_0(s_A)} \geqslant z$: this ensures that the receiver will be willing to adopt after seeing a realization of s_A, and since $\mu < z$ by assumption, this in turn implies that the receiver will be willing to reject after seeing a realization of s_0. For any signal that satisfies this condition, we will have $p_1 = \pi_1(s_A)$, $p_0 = \pi_0(s_R)$, by definition of incentive compatibility. This observation allows us to give a succinct definition of the sender's choice set: $P(\Pi^*) = \{(p_0, p_1) \in [0,1] \times [0,1] : \frac{\mu p_1}{\mu p_1 + (1-\mu)(1-p_0)} \geqslant z\}$. This unrestricted choice set $P(\Pi^*)$ will be useful as a comparison when we restrict the sender's choice set through either pre-registration or sequential sampling.

The sender's optimal trial under the unrestricted choice set will recommend adoption with probability 1 in the good state and probability $\frac{\mu(1-z)}{(1-\mu)z}$ in the bad

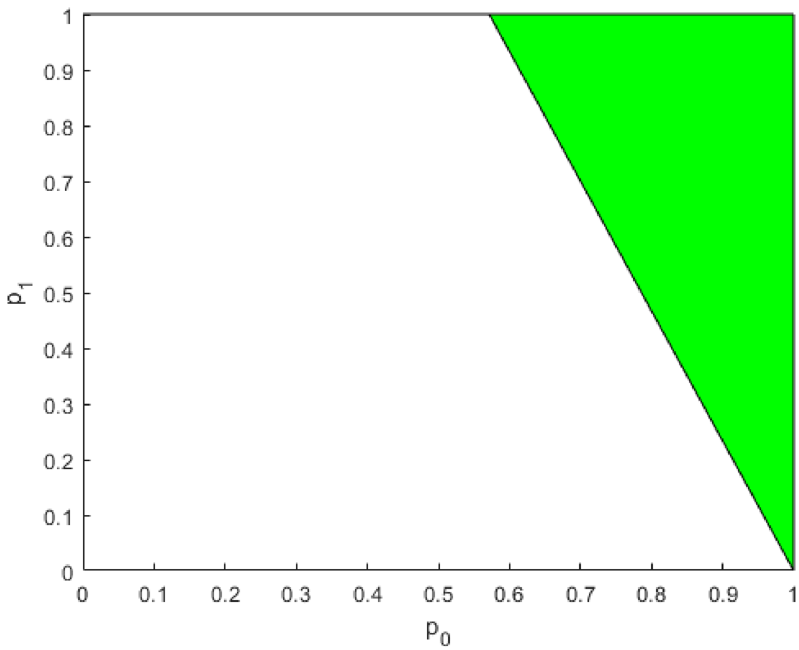

Figure 1: Sender's unrestricted choice set when $\mu = .3, z = .4$. The incentive compatibility constraint that defines the set is $\frac{(.3)p_1}{((.3)p_1+(.7)(1-p_0))} > .4$

state. Thus when the trial recommends adoption, the receiver will believe the probability that the state is good is exactly z. The sender's value from this optimal signal is $\mu + (1-\mu)\frac{\mu(1-z)}{(1-\mu)z} = \mu(1+\frac{(1-z)}{z})$.

3.2 Distribution of Subject Outcomes

In both the pre-registration and sequential sampling regimes, the sender's trial will involve enrolling some number of subjects. To fix ideas, I suppose that subjects respond to treatment in the following symmetric fashion. After receiving treatment, a subject's condition will improve with probability ρ when the state is good, and $1-\rho$ when the state is bad. Letting s_ω designate the outcome that is more common in state ω, we have the following conditional distribution.

	$s = s_0$	$s = s_1$
$\omega = 0$	ρ	$1-\rho$
$\omega = 1$	$1-\rho$	ρ

3.3 The Pre-Registration Choice Set

Under pre-registration, the sender can only choose the sample size of his trial. When the sender chooses sample size n, there are $n+1$ possible payoff-relevant outcomes: the number of subjects that improve after treatment is given by $x \in \{0, 1, ..., n\}$. The probability that x subjects improve when the sample size is $n > x$ is given by $\binom{n}{x}\rho^x(1-\rho)^{n-x}$ in state 1, and $\binom{n}{x}\rho^{n-x}(1-\rho)^x$ in state 0.

Thus for an appropriate choice of n, under pre-registration the sender can choose any information structure in $\Pi^{PR} = \{\pi = (S, \pi_0, \pi_1) \in \Pi^* : \exists n \in \mathbf{N} \text{ s.t. } S = \{0, 1, .., n\}, \text{ and } \forall s \in S, \pi_0(s) = \binom{n}{s}\rho^{n-s}(1-\rho)^s, \pi_1(s) = \binom{n}{s}\rho^s(1-\rho)^{n-s}\}$. Taking $\tilde{P}(\Pi^{PR})$ to be the closure of $P(\Pi^{PR})$, and allowing the sender to randomize over his choice of sample size, I will take the convex hull of $\tilde{P}(\Pi^{PR})$ to be the sender's choice set under pre-registration, which I will refer to as $\bar{P}(\Pi^{PR})$.

Below I have plotted the action distributions induced by values of $n \in \{0, 1, ..., 55\}$, when $\mu = .3$, $z = .5$, and $\rho = .68$. The induced distributions (p_0, p_1) are represented by the blue stars, and the convex hull of the set is shaded in red. The yellow line is the sender's indifference curve associated with full information revelation ($n \to \infty$), and the orange line is the receiver's indifference curve associated with no information revelation ($n = 0$). Note that the unrestricted choice set $P(\Pi^*)$ would be the entire area to the northeast of the orange line.

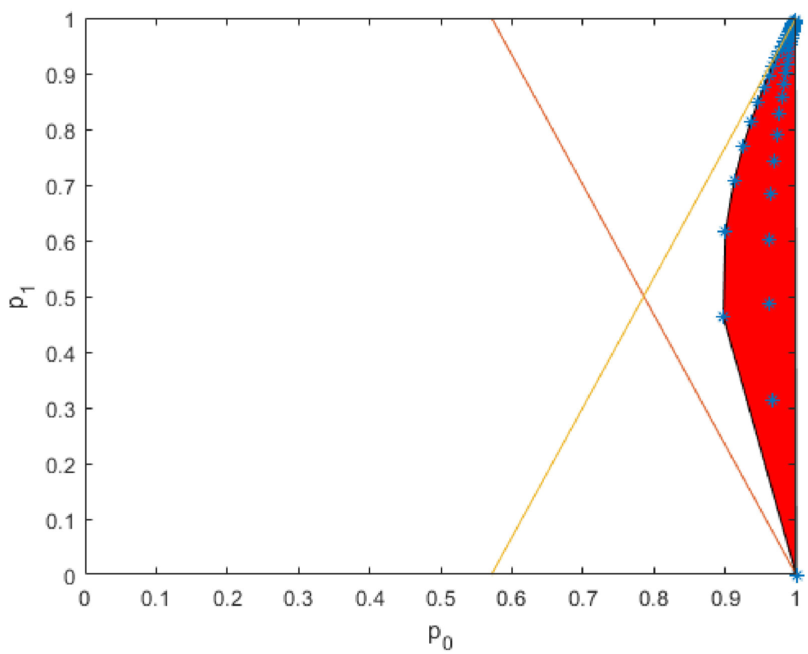

Below is the sender's choice set for the same values of μ and z, where $\rho = .71$.

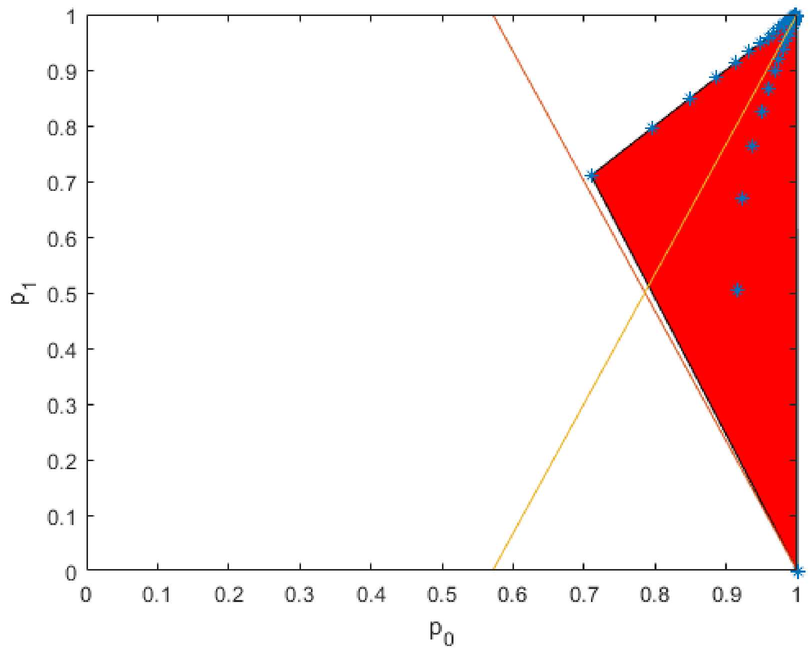

3.4 The Sequential Sampling Choice Set

Under sequential sampling, the sender's problem is simply to choose when to stop enrolling subjects. It is without loss of generality to assume that the sender can commit to a stopping rule (or a randomization over stopping rules) ex ante. A given stopping rule T will induce some distribution over receiver posterior beliefs. Let π^T denote the signal which induces the same distribution over receiver posteriors (and therefore the same distribution over receiver actions) as the stopping rule T. For a more detailed treatment of the static signals induced by stopping rules, see Morris and Strack (2019). Thus we can think of the sender as choosing a signal from $\Pi^{SS} = \{\pi^T : T \text{ is a stopping rule}\}$. This is equivalent to choosing an action distribution from $P(\Pi^{SS})$; as under pre-registration, assume the sender can randomize over stopping rules and so can choose any action distribution from the closed convex hull of $P(\Pi^{SS})$, which I will refer to as $\bar{P}(\Pi^{SS})$.

4 Analysis

Throughout this analysis, I will use "for all μ" to refer to all $\mu \in (0, 1)$, "for all z" to refer to all $z \in (\mu, 1)$, and "for all ρ" to refer to all $\rho \in (.5, 1)$.

A useful feature of the binary symmetric distribution is that that, whatever prior belief μ the observer holds, after seeing one positive and one negative realization of i.i.d. binary symmetric signals, the observer will be back to holding belief μ. Thus we can restrict attention to the difference between successes and failures, which I will define as $d = x - (n - x)$. In the lemma below, I compute the minimum difference d^* that would cause the receiver to adopt. Notably, d^* depends on ρ, z, and μ, but not n. Throughout this paper I will use $\log(\cdot)$ to refer to the natural logarithm, with base e.

Lemma 4.1 *The receiver will adopt if and only if she sees a difference* $d \geqslant d^* = \left\lceil \frac{\log(\frac{1-\mu}{\mu}\frac{z}{1-z})}{\log(\frac{\rho}{1-\rho})} \right\rceil$

Proof The receiver's belief after seeing a trial with a resulting difference of d would be the same as her belief after seeing d positive outcomes in a row. We can write $Pr(\omega = 1|d) = \frac{\mu\rho^d}{\mu\rho^d+(1-\mu)(1-\rho)^d}$. The receiver will thus be willing to adopt in this case if $\frac{\mu\rho^d}{\mu\rho^d+(1-\mu)(1-\rho)^d} > z$. We can simplify this to $\mu\rho^d > z(\mu\rho^d + (1-\mu)(1-\rho)^d)$, further to $(1-z)\rho^d > \frac{1-\mu}{\mu}z(1-\rho)^d$, further still to $(\frac{\rho}{1-\rho})^d > \frac{1-\mu}{\mu}\frac{z}{1-z}$, and finally to $d > \frac{\log(\frac{1-\mu}{\mu}\frac{z}{1-z})}{\log(\frac{\rho}{1-\rho})}$. Take d^* to be the ceiling of this expression, and the proof is complete. ∎

In the next proposition, I show that if the sender's choice set contains his first-best point from the unrestricted choice set $P(\Pi^*)$, then his choice set must be equal to $P(\Pi^*)$.The following proposition has a geometric intuition. The sender's unrestricted Bayesian persuasion choice set is a triangle in the (p_0, p_1) plane, whose three corners are (i) the sender's first-best point, (ii) full revelation, and (iii) no revelation. Since the points (ii) and (iii) are part of the sender's choice set under either regime, if the sender's choice set contains point (i) as well, then by randomizing between the three points, the sender can induce any action distribution in $P(\Pi^*)$.

Proposition 4.2 *Under either regime, if the sender is able to attain the same payoff as he can from* Π^**, then his choice set must be equal to* $P(\Pi^*)$*.*

Proof Recall that $\bar{P}(\Pi^{SS})$ and $\bar{P}(\Pi^{PR})$ are both closed and convex, and both are subsets of $P(\Pi^*)$. Both contain the points (1,0) (achieved when the sender does not enroll any subjects)[3] and (1,1) (achieved as the number of subjects approaches infinity). If the sender can attain his first-best payoff, the upper-left hand corner of $P(\Pi^*)$ must also be included, which implies that the sender's choice set includes all three corners of $P(\Pi^*)$, and therefore must be equal to $P(\Pi^*)$ by convexity. ∎

This shows that the utility function I have assumed for the sender is extreme in the following sense: a sender with utility $v = a$ can attain his first best utility if and only if, for every other function $\hat{v}(a, \omega)$, he would be able to attain his first best utility if his utility function were $\hat{v}(\omega, a)$, as well.

[3]Note that this is due to the assumption $z > \mu$.

4.1 Sequential Sampling

Under sequential sampling, the sender views each subject outcome before deciding whether to enroll another subject, or end the trial. The sender's optimal stopping rule is to stop enrolling subjects if and only if $d \geqslant d^*$. This will convince the receiver to adopt with probability 1 in state 1, since if the sender does not stop producing information, the true state will eventually be revealed. In state 0, the question is a little more complicated, and the problem ultimately reduces to an example of the gambler's ruin problem.

To facilitate the comparison, let D^* denote the event that $d \geqslant d^*$. Let $q_k = Pr(D^{*c}|\omega = 0|d = d^* - k)$ be the probability that D^* is never reached when $\omega = 0$, starting from an observed difference of $d = d^* - k$. We can condition q_k on the value of the next realization, so that $q_k = \rho q_{k+1} + (1-\rho)q_{k-1}$. This can also be written as $\rho q_k + (1-\rho)q_k = \rho q_{k+1} + (1-\rho)q_{k-1}$; from this we can derive $q_{k+1} - q_k = \frac{1-\rho}{\rho}(q_k - q_{k-1})$. Notably, this applies when $k = 1$, yielding $q_2 - q_1 = \frac{1-\rho}{\rho}q_1$, since $q_0 = 0$ by definition. Applying the relationship repeatedly, obtain $q_{k+1} - q_k = (\frac{1-\rho}{\rho})^k q_1$. Then $q_{k+1} - q_1 = \sum_{i=1}^{k} = q_{i+1} - q_i = \sum_{i=1}^{k}(\frac{1-\rho}{\rho})^i q_1$, so $q_{k+1} = \sum_{i=0}^{k}(\frac{1-\rho}{\rho})^i q_1$. Applying our knowledge of geometric series, get $q_{k+1} = q_1 \frac{1-(\frac{1-\rho}{\rho})^{k+1}}{1-\frac{1-\rho}{\rho}}$.

All that is left is to solve for q_1. Assume first that there is some large negative difference $d = d^* - K$ upon reaching which the sender would give up and stop enrolling subjects. Then it would be true that $q_K = 1$. We would also have $1 = q_K = q_1 \frac{1-(\frac{1-\rho}{\rho})^K}{1-\frac{1-\rho}{\rho}}$. This tells us that $q_1 = \frac{1-\frac{1-\rho}{\rho}}{1-(\frac{1-\rho}{\rho})^K}$, so $q^k = \frac{1-(\frac{1-\rho}{\rho})^k}{1-(\frac{1-\rho}{\rho})^K}$. As we let $K \to \infty$, we get $q_k \to 1 - (\frac{1-\rho}{\rho})^k$.

From this, we can conclude that the probability of NOT reaching $d \geqslant d^*$ when $\omega = 0$ and starting with an initial difference of 0 is given by $q_{d^*} = 1 - (\frac{1-\rho}{\rho})^{d^*}$. Thus when $\omega = 0$ the probability of successfully persuading the receiver using sequential sampling is $(\frac{1-\rho}{\rho})^{d^*}$.

Thus the value of sequential sampling to the sender is given by $V^w(\rho) = \mu + (1-\mu)(\frac{1-\rho}{\rho})^{\left\lceil \frac{\log(\frac{1-\mu}{\mu}\frac{z}{1-z})}{\log(\frac{\rho}{1-\rho})} \right\rceil}$.

We are interested in the limit as $\rho \to_+ .5$. Note that in the limit, $d^* = \left\lceil \frac{\log(\frac{z}{1-z})}{\log(\frac{\rho}{1-\rho})} \right\rceil \to_+ \infty$, and so I will drop the ceiling term in what follows.

Write $\lim_{\rho \to_+ .5} \mu + (1-\mu)(\frac{1-\rho}{\rho})^{\frac{\log(\frac{1-\mu}{\mu}\frac{z}{1-z})}{\log(\frac{\rho}{1-\rho})}} = \mu + (1-\mu)(\lim_{\rho \to_+ .5}(\frac{1-\rho}{\rho})^{\frac{\log(\frac{1-\mu}{\mu}\frac{z}{1-z})}{\log(\frac{\rho}{1-\rho})}})$

$$= \mu + (1-\mu)(\lim_{\rho \to_+ .5}[(\frac{1-\rho}{\rho})^{\frac{1}{\log(\frac{\rho}{1-\rho})}}]^{\log(\frac{1-\mu}{\mu}\frac{z}{1-z})}).$$

Note that we can write $(e^y)^{-y^{-1}} = e^{-1}$. Setting $y = \log(\frac{1-\rho}{\rho})$, this identity becomes $(\frac{1-\rho}{\rho})^{\frac{1}{\log(\frac{\rho}{1-\rho})}} = e^{-1}$. Plugging this in, our limit immediately reduces to

$$\mu + (1-\mu)(\lim_{\rho \to_+ .5} [e^{-1}]^{\log(\frac{1-\mu}{\mu}\frac{z}{1-z})}) = \mu + (1-\mu)\frac{\mu}{1-\mu}\frac{1-z}{z} = \mu + \mu(\frac{1-z}{z})$$

which is the same as the sender's first-best payoff under unrestricted Bayesian persuasion. Thus we have the following theorem. Here I use $\pi^*_{SS}(\rho)$ to denote the signal that induces the same distribution over receiver posteriors as the sender's optimal stopping rule, given ρ.

Theorem 4.3 *For all* μ, z, $\lim_{\rho \to_+ .5} v(p_0(\pi^*_{SS}(\rho)), p_1(\pi^*_{SS}(\rho))) = \mu + \mu(\frac{1-z}{z})$ *under sequential sampling.*

This requires that $p(\pi^*_{SS}(\rho))$ converges to the sender's first-best point in Π^*. The following corollary follows.

Corollary 4.4 *Under sequential sampling,* $\lim_{\rho \to_+ .5} u(p_0(\pi^*_{SS}(\rho)), p_1(\pi^*_{SS}(\rho))) = 0$.

4.2 Pre-Registration

Under pre-registration, the sender commits to a non-negative integer sample size n before viewing any subject outcomes, and has the ability to randomize over sample sizes. Since every randomized trial design is weakly dominated by a trial design without randomization, the sender can restrict his attention to signals associated with integer choices of n. Accordingly, for the remainder of this section I will refer to the sender's choice of signal π interchangeably with his choice of sample size n.

Unlike under sequential sampling, calculating the value of an optimal signal to the sender under pre-registration is more difficult than it may seem. The sender's expected utility given a sample size of n is

$$v(p(n)) = \mu p_1(n) + (1-\mu)(1-p_0(n))$$
$$= \mu \sum_{x=x^*}^{n} \binom{n}{x} \rho^x (1-\rho)^{n-x} + (1-\mu) \sum_{x=x^*}^{n} \binom{n}{x} \rho^{n-x} (1-\rho)^*,$$

where again $x^* = \left\lceil \frac{d^*+n}{2} \right\rceil$ and $d^* = \frac{\log(\frac{z(1-\mu)}{\mu(1-z)})}{\log(\frac{\rho}{1-\rho})}$. Consider the abstracted extreme case where $\rho = .5$ and x^* is finite.[4] Then the sender's value becomes $v(n) = (.5)^n \sum_{x=x^*}^{n} \binom{n}{x}$. There is no known closed form expression for the partial sum of binomial coefficients[5]. Here the sender's problem is slightly different than just a sum of binomial coefficients, and there are some methods for evaluating other hypergeometric partial sums (e.g. Petkovsek, Wilf, and Zeilberger 1996), so there may exist a closed form expression for the sender's value. For now, though, I will focus on approaches which do not require us to evaluate the sender's value function.

[4] I say this example is "abstracted" because, by definition, as $\rho \to .5$, $x^* \to \infty$.

[5] "the indefinite sums $\sum_{k=0}^{K_0} \binom{n}{k}$ cannot be expressed in simple hypergeometric terms in K_0 (and n)" (Petkovsek, Wilf, and Zeilberger 1996).

Lemma 4.5 *For all μ, z, ρ, n, $p_0(n) \geqslant p_1(n)$.*

Proof Fix n, μ, z, ρ. We can write

$$p_1(n) = Pr(d \geqslant d^*|n, \omega = 1),$$

$$p_0(n) = Pr(d \leqslant -d^*|n, \omega = 0) + Pr(d^* > d > -d^*|n = n, \omega = 0).$$

By symmetry, $Pr(d \geqslant d^*|n = n, \omega = 1) = Pr(d \leqslant -d^*|n = n, \omega = 0)$. Thus $p_0(n) \geqslant p_1(n)$ for arbitrary μ, z, ρ, n. ∎

Proposition 4.6 *If $\mu > .5$, then for all z and ρ, full revelation (the limit as $n \to \infty$) is the sender's unique optimal choice under pre-registration.*

Proof The sender's payoff is the probability of receiver adoption. Write $v(p) = \mu p_1 + (1 - \mu)(1 - p_0)$. Consider the following maximization problem:

$$\max_{(p_0,p_1)\in[0,1]\times[0,1]} \mu p_1 + (1 - \mu)(1 - p_0)$$

$$s.t. \quad p_0 - p_1 \geqslant 0$$

and observe that as long as $\mu > 1 - \mu$ (i.e. $\mu > .5$), the unique solution is $(p_0, p_1) = (1, 1)$. Since the sender's choice set $P(\bar{\Pi}^{PR})$ is a subset of the choice set here (by the previous lemma), and $(1, 1) \in P(\bar{\Pi}^{PR})$ (full revelation is attained in the limit as $n \to \infty$, and $P(\bar{\Pi}^{PR})$ is closed), this implies that under pre-registration, full revelation is the sender's unique optimal choice when $\mu > .5$. ∎

The following proposition provides a bound on p_0.

Proposition 4.7 *$p_0(n) \geqslant .5$ for all μ, z, ρ, n.*

Proof Fix n, μ, z, ρ. Since $z > \mu$ by assumption, we must have $d^* \geqslant 1$. That is, the receiver needs to see at least one more positive outcome than negative outcome, in order to be willing to adopt. Treating subjects as distinguishable, there are 2^n possible histories of subject outcomes (n subjects, each has one of two outcomes). For every history that results in the receiver choosing to adopt, there is a "mirrored" history, where each subject's outcome is reversed. Such a mirrored history would lead the receiver to reject the sender's treatment. When the state of the world is 0, each history that leads the receiver to adopt is strictly less likely to occur than the corresponding mirrored history (since negative outcomes are more likely than positive outcomes in state 0). Thus, in state 0, the probability that the receiver does not adopt is less than the probability that the receiver adopts. Hence $p_0(n) \geqslant .5$.

Note that the proofs thus far can be applied generally whenever the distribution of subject outcomes is symmetric, e.g. if the distribution of subject outcomes is Normal with mean ω and variance σ^2.

The following proposition uses the additional assumption that subject outcomes are Bernoulli in order to prove a tighter bound on p_0.

Proposition 4.8 *For all* $\mu, z, \rho, n,\ p_0(n) \geqslant 1 - \frac{1}{1+\frac{z(1-\mu)}{\mu(1-z)}}$.

Proof For any history $h = (s_1, s_2, ..., s_n)$ of subject outcomes that would cause the receiver to adopt the sender's product, there is a "mirrored" history, in which all subject outcomes are reversed, after seeing which the receiver would reject the sender's product. Consider a history h' which is identical to h except that one subject's good outcome in h is replaced with one bad outcome for that subject in h'. When $\omega = 0$, history h' is more likely than history h by a factor of $\frac{\rho}{1-\rho}$. This is because in the bad state, a bad subject outcome is more likely than a good subject outcome by a factor of $\frac{\rho}{1-\rho}$. Similarly, replacing a bad outcome with a good outcome multiplies the likelihood of the history by a factor of $\frac{1-\rho}{\rho}$. Replacing one good outcome with a bad outcome and one bad outcome with a good outcome does not affect the likelihood of the history. Since any history h which leads the receiver to adopt must have at least d^* more good outcome than bad outcome, the mirrored history which reverses each subject outcome must be more likely than history h by a factor of at least $(\frac{\rho}{1-\rho})^{d^*}$.

This yields the inequality $p_0 \geqslant (\frac{\rho}{1-\rho})^{d^*}(1 - p_0)$. We can then re-write $1 = (1 - p_0) + p_0$ as $1 \geqslant (1 - p_0) + (\frac{\rho}{1-\rho})^{d^*}(1 - p_0)$. Solving for $(1 - p_0)$, get $(1 - p_0) \leqslant \frac{1}{1+(\frac{\rho}{1-\rho})^{d^*}}$, and so $p_0 \geqslant 1 - \frac{1}{1+(\frac{\rho}{1-\rho})^{d^*}}$.

Note that lowering d^* leads to a smaller quantity being subtracted, which loosens the bound, and recall that $d^* = \left\lceil \frac{\log(\frac{z(1-\mu)}{\mu(1-z)})}{\log(\frac{\rho}{1-\rho})} \right\rceil$. Thus we can write

$$p_0 \geqslant 1 - \frac{1}{1 + (\frac{\rho}{1-\rho})^{\frac{\log(\frac{z(1-\mu)}{\mu(1-z)})}{\log(\frac{\rho}{1-\rho})}}}.$$

Simplifying the denominator, write $(\frac{\rho}{1-\rho})^{\frac{\log(\frac{z(1-\mu)}{\mu(1-z)})}{\log(\frac{\rho}{1-\rho})}} = ((\frac{\rho}{1-\rho})^{\frac{1}{\log(\frac{\rho}{1-\rho})}})^{\log(\frac{z(1-\mu)}{\mu(1-z)})} = e^{\log(\frac{z(1-\mu)}{\mu(1-z)})} = \frac{z(1-\mu)}{\mu(1-z)}$, where the second to last inequality is due to the identity $(e^y)^{(1/y)} = e$, with $y = \log(\frac{\rho}{1-\rho})$. Thus we have the bound $p_0 \geqslant 1 - \frac{1}{1+\frac{z(1-\mu)}{\mu(1-z)}}$.

∎

5 Conclusion

We have seen that in a simple model of trial design, requiring the sender to commit to a sample size before seeing any subject outcomes can lead to a large increase in receiver welfare. Under the sequential smapling regime, as the information generated by each subject vanishes, the sender can approach his first-best Bayesian persuasion payoff, and the receiver's payoff approaches her first-worst. Under the pre-registration regime, we have seen that there is a bound on the false positive rate of the sender's test, and if the good state is at least as likely as the bad state ex ante, the sender will choose to reveal the state fully

Further questions remained to be studied. While I have focused the analysis of sequential sampling on the case when $\rho \to_+ .5$, the case where $\rho > .5$ is fixed is also interesting. It seems likely that there are some values of μ, z, ρ for which the receiver prefers the sequential sampling regime to the pre-registration regime. It would also be interesting to see a more general model where the state space need not be binary, and subject outcomes need not be Bernoulli. That said, I think there are still more results to be proven

6 References

Henry, Emeric, and Marco Ottaviani. 2019. "Research and the Approval Process: The Organization of Persuasion." American Economic Review, 109 (3): 911-55.

Kamenica, E., Gentzkow, M.: Bayesian persuasion. Am. Econ. Rev. 101, 2590–2615 (2011)

Kaplan, R. & Irvin, V. PLoS ONE 10, e0132382 (2015)

Morris, Stephen, and Philipp Strack. 2019. "The Wald Problem and the Equivalence of Sequential Sampling and Static Information Costs." Unpublished.

Petkovšek, Marko., Herbert S. Wilf, and Doron. Zeilberger. A=B / Marko Petkovšek, Herbert S. Wilf, Doron Zeilberger. Wellesley, MA: A K Peters, Ltd., 1996. Print.