Persuasion through Trial Design: Pre-Registration Versus Sequential Sampling

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 $\rho \to_+ .5$, the researcher can achieve his first-best Bayesian persuasion outcome, which minimizes policymaker welfare over attainable BP outcomes. Under pre-registration, as $\rho \to_+ .5$, the researcher's optimal trial approaches full revelation, which maximizes policymaker welfare. I also show that under pre-registration, regardless of ρ, the sender is bounded away from his first-best outcome, and when the good state is at least as likely as the bad state ex ante, full revelation is optimal for the sender. However, when the bad state is more likely than the good state, and subject outcomes are very informative, policymaker welfare is higher under sequential sampling than pre-registration.

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 researcher commits to a sample size before any results are realized. In the *sequential sampling* regime, the researcher observes the results of each subject before deciding whether to enroll the next one, and the trial ends when the researcher 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 researcher seeks to maximize the probability of policymaker adoption.

I compare the researcher's optimal trial under the two regimes to the researcher's optimal trial under unrestricted Bayesian persuasion, where he can design any experiment with any outcome space and any probability distributions that he chooses. The set of attainable outcomes under pre-registration is a subset of the set of attainable outcomes under sequential sampling, which in turn is a subset of the set of attainable Bayesian persuasion outcomes. I show that under sequential sampling, as $\rho \to_+ .5$ the researcher's payoff approaches his first-best Bayesian persuasion payoff, and the policymaker's payoff approaches her first-worst Bayesian persuasion payoff. Under pre-registration, in contrast, I show that as $\rho \to_+ .5$ the researcher's optimal trial approaches full revelation, which yields the policymaker her first-best payoff.

I also show that under pre-registration, regardless of the value of ρ, the researcher is bounded away from his first-best Bayesian persuasion payoff, and when $\mu \geqslant .5$, full revelation is optimal for the researcher. However, this does not imply that pre-registration is always better for policymaker welfare. When $\mu < .5$, and subject outcomes are sufficiently informative, I show that policymaker welfare is higher under sequential sampling than under pre-registration.

These results 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 in a binary-state environment. The second result suggests that it may be much less useful for making predictions about trial design under pre-registration. These results also suggest that when subject outcomes are very uninformative or when $\mu \geqslant .5$, pre-registration leads to more informative trials, while sequential sampling leads to more informative trials when $\mu < .5$ and subject outcomes are very informative.

For a graphical example, suppose $\mu = .3$ and $z = .5$. Any trial the researcher conducts can be parameterized by two values: the probability that it leads the policymaker to adopt the researcher's treatment when it is good, p_1, and the probability that it leads the policymaker to reject the researcher's treatment when it is bad, p_0. If the researcher enrolls no subjects (and thus the policymaker receives no information), the policymaker will always reject the researcher's treatment (since $\mu < z$). This trial is associated with the point (1,0) in the figure below. The policymaker's indifference curve which contains this outcome is the orange line. Under the standard Bayesian persuasion framework, the researcher would be able to induce any outcome to the northeast of that orange indifference curve. Under sequential sampling, as $\rho \to_+ .5$, I show that the researcher can similarly induce any outcome to the northeast of that orange line, just as in the standard BP framework. The purple line graphs $p_0 = p_1$, and I prove that under pre-registration the researcher can not obtain any outcome to the northwest of this purple line, regardless of the value of ρ. In section 4 I show that this bound is tight, in that there is a value of ρ for which the researcher can obtain any outcome that lies above the orange line and below the purple line. The blue stars represent the outcomes the researcher can induce under pre-registration when $\rho = .68$, and by randomizing over his choice of sample size he can obtain any outcome in their convex hull, which is shaded in red. The yellow line here represents the researcher's indifference curve associated with full revelation.

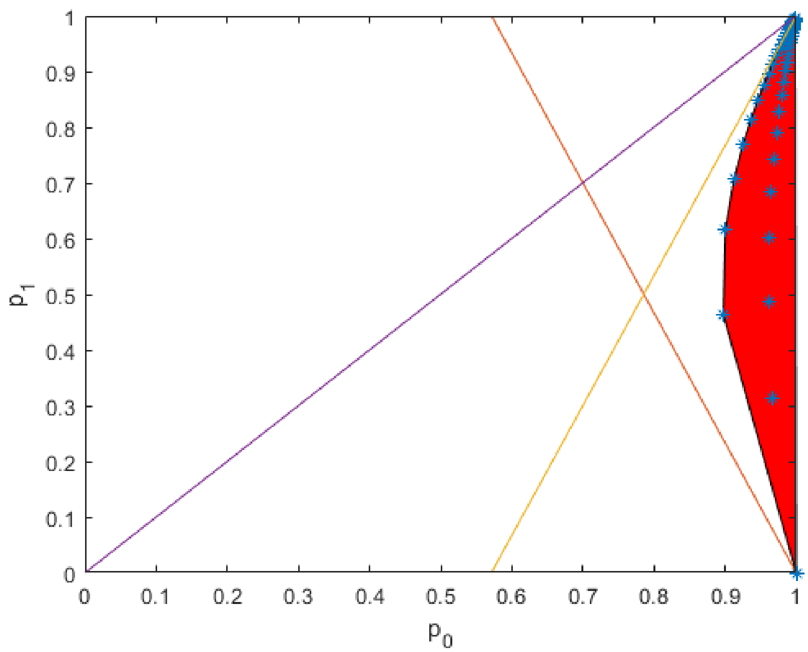

Registration of clinical trials is already the norm in many instances. The International Committee of Medical Journal Editors summarize their policy as follows: "Briefly, the ICMJE requires, and recommends that all medical journal editors require, registration of clinical trials in a public trials registry 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 experimenter bias are more common in fields like psychology without such commitments to pre-registration. Here, "registration" is

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

done through websites such as clinicaltrials.gov, which requires researchers to register their sampling method and target number of participants, 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 from an exogenously given choice set Π, which stochastically maps the state of the world into some outcome space. 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 the prior μ. Given these assumptions, her utility can be normalized as

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

$u(0, \omega) = 0, u(1, \omega) = \omega - z$ for $\omega \in \{0, 1\}$, for some $z \in (\mu, 1)$. With this parameterization, the receiver will strictly prefer to adopt when her posterior that the state is good is greater than z, and prefer to reject when her posterior is less than z. When the receiver's posterior belief is exactly z, I will assume she chooses to adopt the sender' The sender's utility depends only on whether the receiver adopts the sender's treatment, $v = a$.

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 $z > \mu$ by assumption). 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 instead consider his choice to be over action distributions from $P(\Pi)$.

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

As a benchmark, 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 (since the value of rejecting is always 0). 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

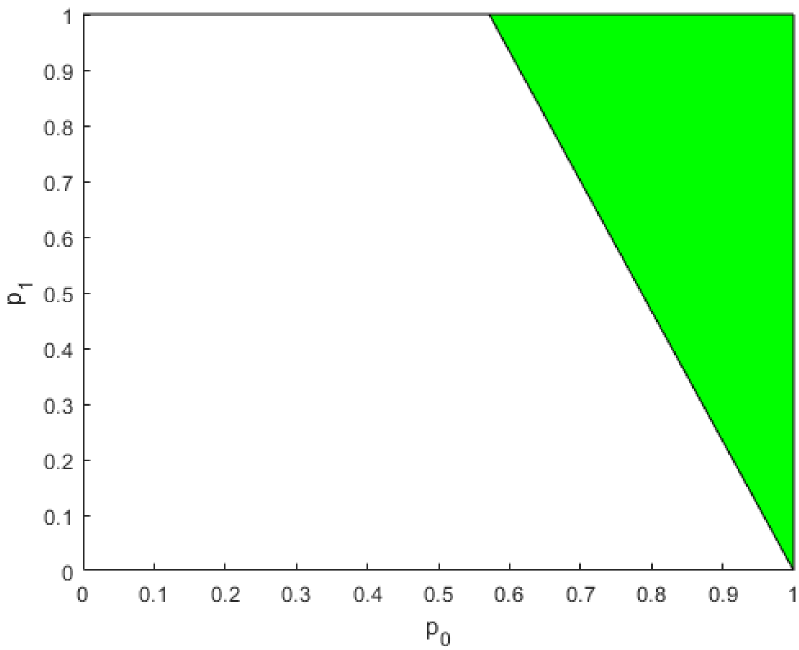

Figure 1: Sender's unrestricted choice set when $\mu = .3, z = .4$. The incentive compatibility constraint that defines the set is $\frac{(.3)p_1}{((.3)p_1+(.7)(1-p_0))} > .4$

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 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 $\mu + (1-\mu)\frac{\mu(1-z)}{(1-\mu)z} = \mu + \mu\frac{(1-z)}{z}$).

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 $\Pi^{PR} = \{\pi = (S, \pi_0, \pi_1) \in \Pi^* : \exists n \in \mathbf{N}$ s.t. $S = \{0, 1, .., n\}$, and $\forall s \in S, \pi_0(s) = \binom{n}{s}\rho^{n-s}(1-\rho)^s, \pi_1(s) = \binom{n}{s}\rho^s(1-\rho)^{n-s}\}$.

Let $\tilde{P}(\Pi^{PR})$ be the closure of $P(\Pi^{PR})$ (so that full revelation is an option for the sender). By randomizing over his choice of sample size, the sender can obtain any convex combination of outcomes in $\tilde{P}(\Pi^{PR})$, so I define the sender's choice set under pre-registration $\bar{P}(\Pi^{PR})$ to be the convex hull of $\tilde{P}(\Pi^{PR})$.

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.

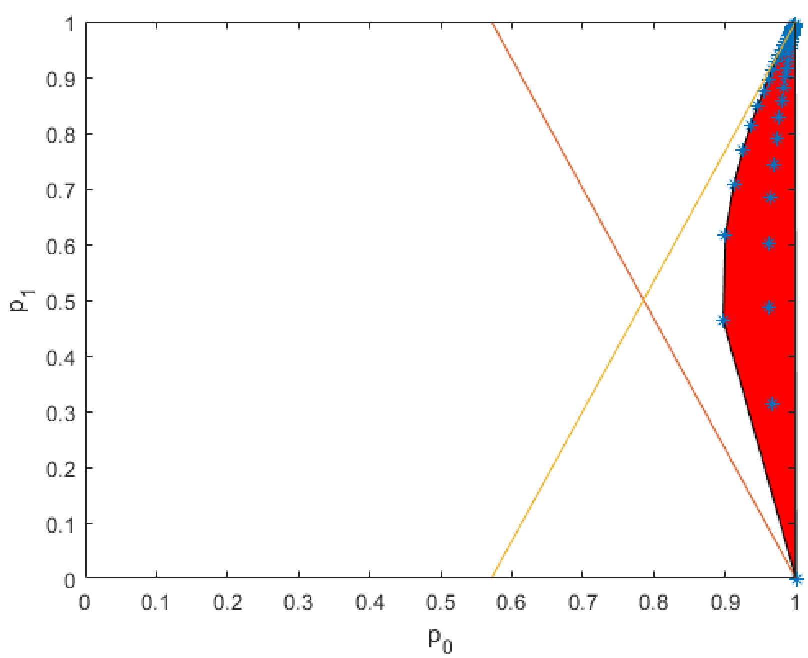

Below is the sender's choice set for the same values of μ and z, where $\rho = .71$.

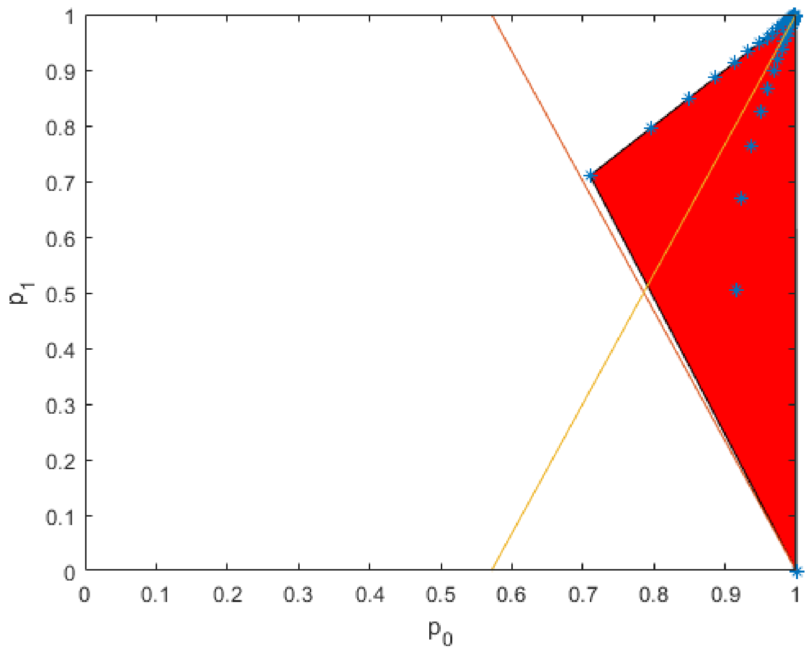

3.4 The Sequential Sampling Choice Set

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.

Formally, define $H_n = \times_{i=1}^{n}\{s_0, s_1\}$ to be the set of all possible histories after seeing n subjects, and $H = \bigcup_n H_n$ to be the set of all possible histories. A stopping rule $T \subset H$ can be defined by the subset of histories at which the sender will choose not to enroll any further subjects.

Any 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 our binary symmetric distribution of subject outcomes is, 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.

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

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 $d > \frac{\ln(\frac{1-\mu}{\mu}\frac{z}{1-z})}{\ln(\frac{\rho}{1-\rho})}$. 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 $P(\Pi^*)$. This 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 $P(\Pi^*)$. Since the sender's choice set is convex, it must be equal to $P(\Pi^*)$. ∎

[3]due to the assumption $z > \mu$.

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.

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. As demonstrated by Brocas and Carillo (2007), the sender's value can be computed using a second-order difference equation. Defining $\hat{z}(d^*) = \frac{\mu\rho^{d^*}}{\mu\rho^{d^*}+(1-\mu)(1-\rho)^{d^*}}$, we can borrow the following proposition.

Proposition 4.3 *(Brocas and Carrillo 2007) Under sequential sampling, the sender's value from following his optimal stopping rule T^* is $v(p(T^*)) = \mu + \frac{\mu}{\hat{z}(d^*)}(1 - \hat{z}(d^*))$*

As the following lemma shows, as $\rho \rightarrow_+ .5$, $\hat{z}(d^*) \rightarrow z$. As a consequence, the sender's utility from his optimal stopping rule approaches his utility from his first-best Bayesian persuasion outcome.

Lemma 4.4 *As $\rho \rightarrow_+ .5$, $\hat{z}(d^*) \rightarrow z$.*

Proof By definition of the posterior belief $\hat{z}(d^*)$, it is weakly greater than z, but weakly less than the posterior which arises after seeing one positive result starting from a prior belief z. Thus we have $z \leqslant \hat{z}(d^*) \leqslant \frac{z\rho}{z\rho+(1-z)(1-\rho)}$. Taking the limit as $\rho \rightarrow_+ .5$, the inequality becomes $z \leqslant \lim_{\rho\rightarrow_+ .5} \hat{z}(d^*) \leqslant z$. ∎

The following theorem follows immediately.

Theorem 4.5 *For all μ, z, $\lim_{\rho\rightarrow_+ .5} v(p_0(\pi^*_{SS}(\rho)), p_1(\pi^*_{SS}(\rho))) = \mu + \mu(\frac{1-z}{z})$ under sequential sampling.*

This requires that $p(\pi^*_{SS}(\rho))$ converges to the sender's first-best point in Π^*. Since this point lies on the same receiver indifference curve as (1,0), we have the following corollary.

Corollary 4.6 *Under sequential sampling,* $\lim_{\rho\rightarrow_+ .5} u(p_0(\pi^*_{SS}(\rho)), p_1(\pi^*_{SS}(\rho))) = 0$.

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

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

where again $x^* = \left\lceil \frac{d^*+n}{2} \right\rceil$ and $d^* = \frac{\ln(\frac{z(1-\mu)}{\mu(1-z)})}{\ln(\frac{\rho}{1-\rho})}$. Solving for the sender's optimal choice of n is difficult because there is no general hypergeometric expression for the partial sum of binomial coefficients.[4] 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 (see for example 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.

Specifically, I will focus on what happens as $\rho \to_+ .5$. In contrast to the sequential sampling regime, I will show that under pre-registration, regardless of the choice of $n(\rho)$, we have $p_0(n(\rho)) \to 1$ as $\rho \to_+ .5$. The proof of this begins with Hoeffding's inequality, which lets us write $1 - p_0(n) \leqslant e^{-2(\frac{1}{2}+\frac{d^*}{2n}-(1-\rho))^2 n}$. Next, I take the first derivative with respect to n to show that this upper bound is decreasing in n. Finally, I evaluate the bound when $n = d^*$ and show that as $\rho \to_+ .5$, the bound approaches 0. Since the researcher will never choose $n < d^*$, and the bound is decreasing in n, we can conclude that for any choice of $n(\rho)$, we must have $p_0(n(\rho)) \to 1$. This yields the following theorem.

Theorem 4.7 *As* $\rho \to_+ .5, p(\pi^*_{PR}(\rho)) \to (1,1)$.

The full proof can be found in the appendix. This result highlights a sharp contrast between pre-registration and sequential sampling as subject outcomes become uninformative. Under the former, the reseacher's optimal trial converges to full revelation, maximizing policymaker welfare. Under the latter, the researcher's optimal trial converges to his first-best Bayesian persuasion outcome, minimizing policymaker welfare.

[4]"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).

We can also use the symmetry of subject outcomes to derive another bound on the set of inducible outcomes. This bound is "tight" in the sense that there is some value of ρ for which the researcher can attain any Bayesian persuasion outcome that satisfies the bound.

Lemma 4.8 *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. ∎

The following graphs depict this bound, as well as the receiver's indifference curve associated with no revelation, the sender's indifference curve associated with full revelation, and the sender's pre-registration choice set, when $\mu = .3, z = .5$, for $\rho = .68$, $\rho = .71$, and $\rho = .70001$, respectively.

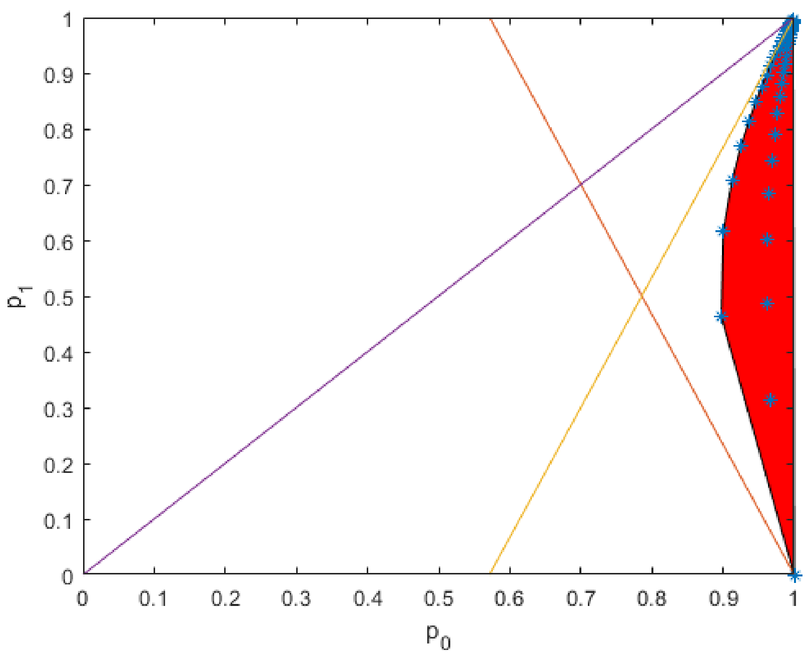

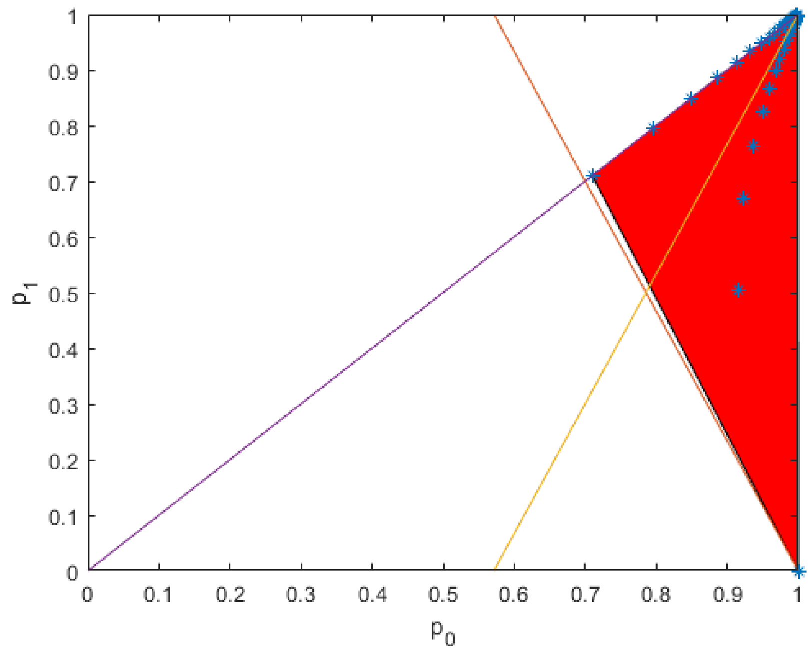

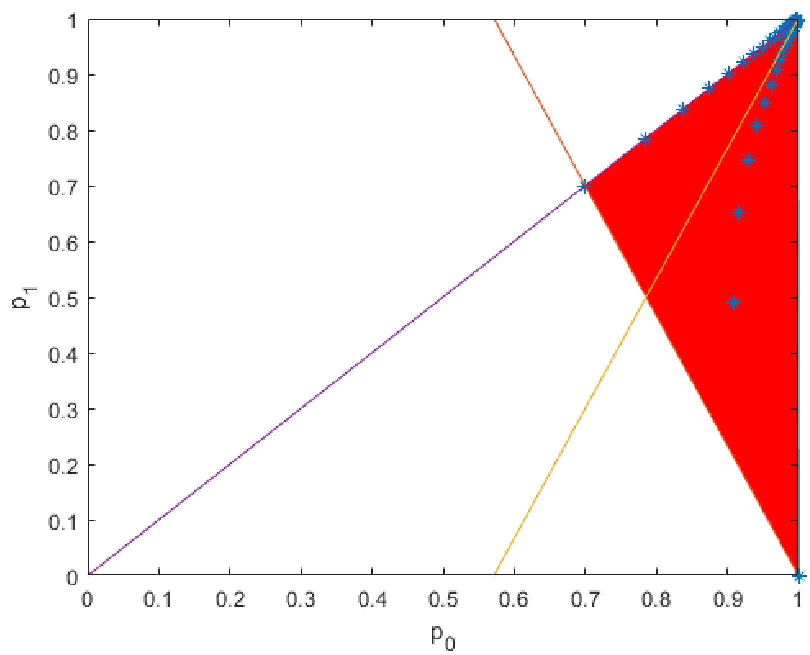

Note that this bound can be attained when $\rho = .7$. At this value, seeing one positive outcome will be exactly informative enough to make the receiver indifferent between adoption and rejection (in which case she will adopt by assumption). This implies that the receiver earns a guaranteed payoff of 0 from seeing a one-subject trial

Thus, there is a sequence of $\{\rho\}$ for which the receiver's utility approaches his baseline utility under pre-registration as well. However, unlike under sequential sampling, where receiver utility vanishes for very small ρ (and thus very large sample sizes), under pre-registration receiver utility vanishes when ρ is so large that the sender's optimal sample only enrolls one subject.

In fact, whenever $\mu < .5$ and ρ is sufficiently large, such a one-subject trial is the sender's optimal trial under pre-registration. To calculate the smallest value of ρ for which the receiver will be willing to adopt after seeing only one positive outcome, I set $z = \frac{\mu\rho}{\mu\rho+(1-\mu)(1-\rho)}$, multiply to obtain $\mu\rho = z(\mu\rho+(1-\mu)(1-\rho))$, combine terms to obtain $(1-z)\mu\rho = z(1-\mu)(1-\rho)$, divide to obtain $\rho = \frac{z(1-\mu)}{(1-z)\mu}(1-\rho)$, add to obtain $\rho(1+\frac{z(1-\mu)}{(1-z)\mu}) = \frac{z(1-\mu)}{(1-z)\mu}$, and divide to obtain $\rho = \frac{\frac{z(1-\mu)}{(1-z)\mu}}{1+\frac{z(1-\mu)}{(1-z)\mu}}$.

Then when $\rho > \frac{\frac{z(1-\mu)}{(1-z)\mu}}{1+\frac{z(1-\mu)}{(1-z)\mu}}$ and the sender only enrolls one subject, the resulting trial will have $p_0 = p_1$, found at the northwestern corner of his feasible set. When $\mu < .5$, the sender's indifference curves are steeper than the bound defined by the line $p_0 = p_1$, and so the sender will prefer this northwestern point to any other point in his feasible set. I formalize this intuition in the following proposition

Proposition 4.9 *When $\mu < .5$ and $\rho > \frac{\frac{z(1-\mu)}{(1-z)\mu}}{1+\frac{z(1-\mu)}{(1-z)\mu}}$, the sender's optimal trial under pre-registration enrolls only one subject.*

Proof When $\rho > \frac{\frac{z(1-\mu)}{(1-z)\mu}}{1+\frac{z(1-\mu)}{(1-z)\mu}}$, if the sender enrolls one subject in his trial, then the receiver will adopt if the outcome is good, and reject if the outcome is bad. As a result, the probability of the receiver choosing correctly is ρ in either state, so $p_0(1) = \rho = p_1(1)$. Note that this lies exactly on the bound defined by the line $p_0 = p_1$. Recall that the sender's expected payoff is $v(p) = \mu p_1 + (1-\mu)(1-p_0)$. Since $p_0(n) \geqslant p_1(n)$, choosing a sample size $n > 1$ can have one of two effects on p_0, p_1: either (i) $p_0(n) = p_1(n)$ or (ii) $p_0(n) \geqslant p_1(n)$. If (i) is true, then adding the additional $n-1$ subjects to the trial has the effect of increasing both p_0 and p_1 by the same amount (and we know they must have increased since the receiver makes more accurate decisions with more information). However, given the sender's expected payoff of $v(p) = \mu p_1 + (1-\mu)(1-p_0)$, when $\mu < .5$, increasing both p_0 and p_1 by the same positive amount can only lower the sender's expected payoff. Then if (ii) is true, we must have p_0 increase by even more than p_1, which similarly can only lower the sender's expected payoff. Thus for any sample size $n \geqslant 1$, the sender's expected payoff is maximized by choosing $n = 1$. ∎

As $\rho \rightarrow_+ \frac{\frac{z(1-\mu)}{(1-z)\mu}}{1+\frac{z(1-\mu)}{(1-z)\mu}}$, the receiver becomes less and less certain about her decision. In the limit where $\rho = \frac{\frac{z(1-\mu)}{(1-z)\mu}}{1+\frac{z(1-\mu)}{(1-z)\mu}}$, upon seeing a good outcome, the receiver is exactly indifferent between adopting and rejecting the sender's treatment. Thus, the receiver is no better off than she was before seeing the sender's trial, when she would choose to reject with probability 1.

Proposition 4.10 *When $\mu < .5$, and $\rho > \frac{\frac{z(1-\mu)}{(1-z)\mu}}{1+\frac{z(1-\mu)}{(1-z)\mu}}$, the receiver's expected utility is higher under sequential sampling than pre-registration.*

Proof When $\mu < .5$, and $\rho > \frac{\frac{z(1-\mu)}{(1-z)\mu}}{1+\frac{z(1-\mu)}{(1-z)\mu}}$, the sender's optimal trial under pre-registration enrolls just one subject. The receiver will adopt if and only if the subject improves. The receiver's expected payoff from adopting is positive, and her payoff from rejecting is 0. Under sequential sampling, the sender begins by enrolling his first subject. There are two cases: either (i) the subject improves, or (ii) the subject does not improve. If (i), the researcher stops the trial, and the receiver adopts. In this case, the receiver's action and expected payoff are exactly the same as under pre-registration. In case (ii), however, the receiver would earn payoff 0 under pre-registration (since she would reject the sender's treatment), but under sequential sampling she earns a positive expected utility (since the sender will continue enrolling subjects, so there is a positive probability that the receiver will adopt the sender's treatment, which gives her positive expected utility).

Note that in this case, pre-registration is worse for the receiver than sequential sampling: under pre-registration she will only see one subject outcome, but under sequential sampling she may see further information.

The flip side of this is that when $\mu > .5$, the sender's indifference curves are less steep than the bound graphed by $p_0 = p_1$, and so the sender strictly prefers the northeastern corner of his feasible set. Thus when $\mu > .5$, full revelation is optimal for the sender. This is formalized in the following proposition.

Proposition 4.11 *If $\mu > .5$, then for all z and $\rho > \rho^*$, 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 Lemma 4.5), 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 $\mu > .5$. ∎

5 Conclusion

We have seen that in a simple model of trial design, requiring the sender to pre-register his sample size can have large effects on receiver welfare and the set of attainable outcomes. Under the sequential sampling 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. At the same time, under the pre-registration regime the sender's optimal trial approaches full revelation, which maximizes receiver payoff. For general ρ, we have also seen that under pre-registration 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. Interestingly, it is not always true that the receiver is better off under pre-registration; when $\mu < .5$ and ρ is sufficiently large, the sender will only pre-register one subject, which is worse for the receiver than sequential sampling. Thus we have the following takeaway: when $\mu < .5$, and when each subject outcome is very uninformative, pre-registration leads to more informative trials, but when subject outcomes are very informative, sequential sampling leads to more informative trials.

Further questions remained to be studied. While the receiver prefers pre-registration when subject outcomes are highly uninformative and sequential sampling when subject outcomes are highly informative, it is still unclear what we can say about intermediate values of ρ. 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.

6 Appendix

Proof of Theorem 4.7 The following proof uses Hoeffding's inequality, which states that for independent random variables $X_1, ..., X_n$, with $X_i \in [0,1]$ for all i,

$$Pr(\sum_i X_i - E[\sum_i X_i] \geqslant t * n) \leqslant e^{-2nt^2}$$

(Hoeffding 1994). For our purposes, each random variable X_i will be the outcome of a subject, such that $X_i = 1$ if subject i improves, and $X_i = 0$ if not. Denote the number of subjects that improve by $x = \sum_i X_i$. When the researcher's treatment is bad ($\omega = 0$), we have $E[X_i] = 1 - \rho$. The policy-maker will adopt the researcher's treatment if $x \geqslant \frac{n+d^*}{2}$, so we are interested in bounding the probability that $x \geqslant \frac{n+d^*}{2}$. Write

$$1 - p_0 = Pr(x \geqslant \frac{n+d^*}{2} | \omega = 0) = Pr(x - n(1-\rho) \geqslant \frac{n+d^*}{2} - n(1-\rho) | \omega = 0)$$

$$= Pr(x - n(1-\rho) \geqslant \frac{\frac{n+d^*}{2} - n(1-\rho)}{n} * n | \omega = 0) \leqslant e^{-2(\frac{1}{2} + \frac{d^*}{2n} - (1-\rho))^2 n}.$$

To show that $1-p_0$ approaches 0 as ρ decreases to .5, I will first show that for any value of ρ, $e^{-2(\frac{1}{2}+\frac{d^*}{2n}-(1-\rho))^2 n}$ is decreasing in n. Then I will show that when $n=d^*$, as $\rho \to_+ .5$, $e^{-2(\frac{1}{2}+\frac{d^*}{2n}-(1-\rho))^2 n} \to 0$.

To show that $e^{-2(\frac{1}{2}+\frac{d^*}{2n}-(1-\rho))^2 n}$ is decreasing in n, note that since $\frac{1}{2} > 1-\rho$ for all $\rho \in (.5, 1)$, the exponent $-2(\frac{1}{2}+\frac{d^*}{2n}-(1-\rho))^2 n$ is always negative. Taking the derivative of the exponent, obtain

$$\frac{d}{dn}[-2(\frac{1}{2}+\frac{d^*}{2n}-(1-\rho))^2 n]$$

$$= -2[2(\frac{1}{2}+\frac{d^*}{2n}-(1-\rho))(-\frac{d^*}{2n^2})n + (\frac{1}{2}+\frac{d^*}{2n}-(1-\rho))^2]$$

$$= -2(\frac{1}{2}+\frac{d^*}{2n}-(1-\rho))[(-\frac{d^*}{2n}) + (\frac{1}{2}+\frac{d^*}{2n}-(1-\rho))]$$

$$= -2(\frac{1}{2}+\frac{d^*}{2n}-(1-\rho))[\frac{1}{2}-(1-\rho)]$$

Since $\frac{1}{2} > 1-\rho$ for all $\rho \in (.5, 1)$, this derivative is negative for all $n > 0$. Thus the exponent is becoming more negative, and the overall bound is becoming tighter, as n increases.

Recall that the researcher will never choose a value of $n < d^*$, since any such choice of n would earn him payoff 0. When $n = d^*$, the bound becomes

$$1 - p_0 \leqslant e^{-2(\frac{1}{2}+\frac{1}{2}-(1-\rho))^2 d^*}$$

As $\rho \to_+ .5$, $d^* \to \infty$, and $e^{-2(\frac{1}{2}+\frac{1}{2}-(1-\rho))^2 d^*} \to 0$. Thus we must have $1 - p_0(d^*) \to 0$, and since $1-p_0(d^*) \geqslant 1-p_0(n) \geqslant 0$ for all n, we have $0 \geqslant \lim_{\rho \to_+ .5}(1-p_0(n)) \geqslant 0$ for all n as well.

7 References

Brocas, Isabelle, and Juan D. Carrillo. "Influence through ignorance." The RAND Journal of Economics 38.4 (2007): 931-947.

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

Hoeffding, Wassily. "Probability inequalities for sums of bounded random variables." The collected works of Wassily Hoeffding. Springer, New York, NY, 1994. 409-426.

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.