A Notation Table

Symbol	Usage
$X \subset \{0,1\}$	Bernoulli random variable indicating
	whether a product was effective on a random instance
q	PDF of the effectiveness distribution
Q	CDF of the effectiveness distribution Q.
μ_0	product effectiveness
μ_b	baseline effectiveness
σ_0	variance of a product's effectiveness: $\sigma_0 = \sqrt{(1-\mu_0)\mu_0}$
σ_b	variance of a baseline product's effectiveness: $\sigma_b = \sqrt{(1-\mu_b)\mu_b}$
c_0	fixed cost of running a clinical trial
c	per-sample cost
R	revenue upon approval
α	p-value threshold used by the principal
$p(\mu_0, n; \mu_b)$	the probability of observing outcomes at least
	as good as the evidence conditioned on the null hypothesis μ_b
$z_{\alpha,n}$	The critical region given n samples and p-value α
$\text{Pass}(\alpha, \mu_0, n)$	the probability that a product with effectiveness μ_0 is approved
$u(\alpha, \mu_0, n)$	utility for an agent with effective size μ_0 and number of samples n
$n_{\mu_0}(\alpha)$	optimal number of samples given p-value threshold α and effectiveness μ_0
$\mu_\tau(\alpha)$	The participation threshold belief under a p-value threshold α
α^*	The critical p-value threshold such that $\mu_\tau(\alpha^*) = \mu_b$

Table 3: Primary Notation

B Omitted proofs in Section 3

Proof of Theorem 3.1

Proof. We first define the following variables for brevity. Let $\Delta_\mu = \mu_0 - \mu_b$ denote the effect size, $\sigma_0 = \sqrt{\mu_0(1-\mu_0)}$ denote the standard deviation of the true process and $\sigma_b = \sqrt{\mu_b(1-\mu_b)}$ the standard deviation of the baseline Bernoulli variable. We use $\phi(z)$ to denote the standard normal, and $\Phi(z)$ to denote its CDF. Lastly, let $d_\alpha = \Phi^{-1}(1-\alpha)$ denote the $1-\alpha$ percentile of the standard normal. Recall that we work under the normal approximation of the binomial distribution where the critical threshold is:

$$z'_{\alpha,n} = \left\{ k \in \mathbb{R} \middle| \int_k^\infty \mathcal{N}(n\mu_b, n\mu_b(1-\mu_b)) \right\} = n\mu_b + \Phi^{-1}(1-\alpha)\sqrt{n\mu_b(1-\mu_b)} = \mu_b n + d_\alpha \sigma_b \sqrt{n} \tag{4}$$

Thus, we may write the pass probability as follows, from which the expected utility expression follows:

$$\begin{aligned} \text{Pass}(\cdot) &= 1 - \Phi\left(\frac{z'_{n,\alpha} - n\mu_0}{\sqrt{n\mu_0(1-\mu_0)}}\right) = 1 - \Phi\left(\frac{n\mu_b + d_\alpha\sqrt{n\mu_b(1-\mu_b)} - n\mu_0}{\sqrt{n\mu_0(1-\mu_0)}}\right) \\ &= 1 - \Phi\left(\frac{d_\alpha\sigma_b}{\sigma_0} - \frac{\Delta_\mu\sqrt{n}}{\sigma_0}\right) \\ \implies u(n;\alpha,\mu_0,\mu_b) &= R - R\Phi\left(\frac{d_\alpha\sigma_b}{\sigma_0} - \frac{\Delta_\mu\sqrt{n}}{\sigma_0}\right) - cn - c_0 \end{aligned}$$

Observe that when the drug is not effective, i.e. $\mu_0 < \mu_b$, then the argument to $\Phi()$ increases in n. Hence, the passing probability, and thus the expected revenue, decreases with increasing n. It is evident that the cost also increases in n. Thus, when the underlying drug is not effective, it is optimal for the agent to choose the smallest value n_{min} possible. Computing the pass probability oracle with n_{min}, they can determine their maximum expected utility and choose to participate if this is positive. In other words, when $\mu_0 < \mu_b$, a constant amount of computation suffices to determine the optimal participation decision.

We next turn to the more interesting case where $\mu_0 \geq \mu_b$. Our goal is to efficiently search the sample space by leveraging structural properties of the utility function. Letting $v = \frac{d_\alpha\sigma_b}{\sigma_0} - \frac{\Delta_\mu\sqrt{n}}{\sigma_0}$, the derivative of the utility function with respect to n is as follows:

$$\frac{\partial u}{\partial n} = -R\frac{\partial\Phi}{\partial v}\frac{\partial v}{\partial n} - c \tag{5}$$

$$= \phi(v)\frac{R\Delta_\mu}{2\sigma_0\sqrt{n}} - c = \frac{R\Delta_\mu}{2\sigma_0\sqrt{n}}\phi\left(\frac{d_\alpha\sigma_b}{\sigma_0} - \frac{\Delta_\mu\sqrt{n}}{\sigma_0}\right) - c \tag{6}$$

The first term is always positive, and since the image of $\phi()$ is always above 0. As n increases, this first term tends to 0, and the derivative is dominated by the second term and becomes negative. In other words, increasing n only increases utility to a point. To precisely understand this characteristic, consider the second derivative of the utility function:

$$\frac{\partial^2 u}{\partial n^2} = \frac{-R\Delta_\mu}{2\sigma_0 n^{3/2}}\phi(v) + \frac{R\Delta_\mu^2}{2\sigma_0^2 n}v\phi(v) = \frac{R\phi(v)\Delta_\mu}{2\sigma_0 n^{3/2}}\left[\frac{\Delta_\mu v}{2\sigma_0}\sqrt{n} - 1\right] \tag{7}$$

The sign of the second derivative, and thus the concavity/convexity properties of the utility function depend solely on $\left[\frac{\Delta_\mu v}{2\sigma_0}\sqrt{n} - 1\right]$ since the term multiplying it is always positive. We now expand it by plugging in the definition of v:

$$\left[\frac{\Delta_\mu v}{2\sigma_0}\sqrt{n} - 1\right] = \frac{-\Delta_\mu}{2\sigma_0^2}n + \frac{d_\alpha\Delta_\mu\sigma_b}{2\sigma_0^2}\sqrt{n} - 1 = \frac{-\Delta_\mu}{2\sigma_0^2}t^2 + \frac{d_\alpha\Delta_\mu\sigma_b}{2\sigma_0^2}t - 1 \tag{8}$$

where we substitute in $t = \sqrt{n}$. Observe that this is a negative quadratic expression with only positive real root being meaningful (since $\sqrt{n}$ cannot be negative). The roots n_1, n_2 can easily be computed by applying the quadratic formula to the instance parameters. We categorize the outcome as follows:

1. No positive roots $\implies$ the function is always *concave*
2. One positive root at $n_1 \implies$ the function is *convex* over $[0, n_1]$ and *concave* over $[n_1, \infty)$.
3. Two positive root at n_1, n_2 with $n_2 > n_1 \implies$ that over $[0, n_1]$ the function is *concave*, over $[n_1, n_2]$ is is *convex* and over $[n_2, \infty)$ it is *concave*.

We can consider the problem of optimizing n over each of these convex/concave regions defined by the roots. Note that a convex function always attains its maximum on the boundary. Thus, for a the convex interval, it suffices to just check the boundary points, of which there are only two. At the optimal n for a concave interval, the first derivative is always 0, or it is a boundary point. Since the latter case is similar to the first one, it suffices to find where the first derivative is 0, if it exists. Since the first derivative is always increasing (monotonic) and we need to find the x-intercept, a binary search suffices: starting in the middle of the interval, increase n if the first derivative is positive, and decrease it if negative. This makes at most $\log n_{max}$ calls. It is immediate that taking max over the best n from each of the at most three yields the globally optimal $n_{\mu_0}(\alpha)$. The agent can compute the corresponding passing probability and utility and decide to participate if this is positive. $\square$

Proof of Lemma 3.1

Proof. We start with the first direction. If an agent with prior belief μ_0 participates with n samples, then by individual rationality, they must have non-negative utility. That is, $\Pr[p(n, \mu_0; \mu_b) \le \alpha] \cdot R \ge (c \cdot n + c_0)$, where $(c \cdot n + c_0)$ is the cost. Now consider an agent with a prior belief $\mu_1 \ge \mu_0$ and using the same number of samples-n. Observe that if $\Pr[p(n, \mu_1; \mu_b)) \le \alpha] \ge Pr[p(n', \mu_0; \mu_b)) \le \alpha]$, then the agent will participate under belief μ_1. To compute $\Pr[p(n, \mu_1; \mu_b) \le \alpha]$, let $n\hat{\mu} = \sum_{i=1}^{n} X_i$ denote the observed outcomes of n samples drawn independently from the distribution with effectiveness μ_1. Observe that since the same number of samples are used here as when the belief was μ_0, the critical region $z_{\alpha,n}$ does not change since it depends only on α, n and μ_b. The probability that these samples, drawn with respect to belief μ_1, will lie in this critical region:

$$\sum_{i \in z_{\alpha,n}} \binom{n}{i} \mu_1^i (1-\mu_1)^{n-i} \ge \sum_{i \in z_{\alpha,n}} \binom{n}{i} \mu_0^i (1-\mu_0)^{n-i}$$

where the inequality follows immediately since $\mu_1 \ge \mu_0$.

For the reverse, consider an agent with effectiveness μ_0' not participating. This means that *for all* n, we have $\Pr[p(n, \mu_0'; \mu_b)) \le \alpha] \cdot R_0 < (c \cdot n + c_0)$. Then for an agent with belief $\mu_1' < \mu_0'$ the following holds:

$$\forall\, n \sum_{i \in z_{\alpha,n}} \binom{n}{i} {\mu'}_1^i (1-{\mu'}_1)^{n-i} < \sum_{i \in z_{\alpha,n}} \binom{n}{i} {\mu'}_0^i (1-{\mu'}_0)^{n-i}$$

In other words, for each n, the passing probability under μ_1' is worse than μ_0', and the agent was already not participating under μ_0' for any n. $\square$

Proof for Lemma 3.2

Proof. We first consider solving for $\mu_\tau(\alpha)$, given an α. To solve this upto some accuracy ε, we can discretize the belief space into $k = \frac{1}{\varepsilon}$ intervals of size ε, and use the mid-point of the interval as its representative belief. The monotonicity of beliefs implies that we can run a binary search over these intervals. Starting with the middle interval $\frac{k}{2}$, if the agent participates at its representative belief, we need not search intervals larger than this. Similarly, if the agent does not participate, we need not search all intervals smaller than this. Clearly, this terminates in $\log\left(\frac{1}{\varepsilon}\right)$ calls to the pass probability oracle. To check participation, we need to compute the optimal n at every belief. Assume the maximum number of sample is n_{max}, it will again take $\log(n_{max})$ to search for the optimal number of samples. So the total complexity should be $\log\left(\frac{1}{\varepsilon}\right) * \log(n_{max})$.

Next, we show that $\mu_\tau(\alpha)$ is non-increasing as α increases. For any α, we know that at the corresponding threshold belief $\mu_\tau(\alpha)$ and its corresponding optimal sample size $n_{\mu_\tau}(\alpha) \triangleq n_\tau(\alpha)$, the utility is 0. In other words: $R \cdot \text{Pass}(n_\tau(\alpha), \mu_\tau(\alpha), \mu_b, \alpha) = cn_\tau(\alpha) + c_0$. It is known that for a fixed effect size and number of samples, the passing probability of a hypothesis test increases in α. In other words, for two p-values α_1, α_2 where $\alpha_1 \leq \alpha_2$, the following holds:

$$\text{Pass}(\alpha_1, \mu_\tau(\alpha_1), n_\tau(\alpha_1)) \leq \text{Pass}(\alpha_2, \mu_\tau(\alpha_2), n_\tau(\alpha_2)) \quad \text{and} \quad u(n_\tau(\alpha_1), \mu_\tau(\alpha_1), \alpha_2) \geq 0$$

Since the agent with effectiveness $\mu_\tau(\alpha_1)$ will have non-negative utility when using $n_\tau(\alpha_1)$ samples when the p-value is α_2, this agent will participate at this p-value (but not necessarily using $n_\tau(\alpha_1)$ samples as that may not be optimal). The monotonicity of participation (Lemma 3.1) implies that under α_2, agents with effectiveness greater than $\mu_\tau(\alpha_1)$ will also participation. Thus, the participation threshold under α_2 by definition must be to the left of $\mu_\tau(\alpha_1)$ - i.e. $\mu_\tau(\alpha_2) \leq \mu_\tau(\alpha_1)$ as desired. □

Proof for Lemma 3.3

Proof. We prove by contradiction. Assume that there exists an α and $\mu_1 \geq \mu_0$ such that $u(\alpha, \mu_0, n_{\mu_0}(\alpha)) \geq u(\alpha, \mu_1, n_{\mu_1}(\alpha))$. First note that if the agent is not participating in either, then the utility is always 0. If they participate in one and not the other, the monotonicity of participation (Lemma 3.1) means it must be under μ_1 (giving positive utility), with μ_b being 0. Thus, the only situation where the initial claim could hold if the agent participates under both. We divide this into the following three cases:

- $\mu_0 \leq \mu_1 \leq \mu_b$: since both μ_0 and μ_1 are less effective than the baseline drug, Theorem 3.1 implies that in both cases, n_{min} samples are used, and in Lemma 3.1 we know that in such settings, for a fixed n, the pass probability increases in μ. Thus the utility under μ_1 cannot be lower than μ_0.
- $\mu_0 \leq \mu_b \leq \mu_1$ and $\mu_b \leq \mu_0 \leq \mu_1$: We know from Lemma 3.1 for every fixed n, the pass probability in this regime increases in μ. Thus, $\forall n, u(\alpha, \mu_0, n) \leq u(\alpha, \mu_1, n)$. Since $n_{\mu_1}(\alpha)$ is the optimal number of samples for μ_1, we have

  $$u(\alpha, \mu_0, n_{\mu_1}(\alpha)) \leq u(\alpha, \mu_1, n_{\mu_1}(\alpha)) \leq u(\alpha, \mu_0, n_{\mu_0}(\alpha))$$

  The last step is according to the contradiction statement. However, this cannot be true since an $n_{\mu_1}(\alpha)$ could not be the optimal number of samples for μ_1 if it led to a lower utility than $n_{\mu_0}(\alpha)$.

□

C Omitted Proof for Section 4

For the ease of notation, let Q denote the CDF of the effectiveness distribution q. Further, let $Q_{<\mu_b}$ and $Q_{\geq\mu_b}$ denote the CDF of the belief distribution q, conditioned on $\mu_0 < \mu_b$ and $\mu_0 \geq \mu_b$.

Proof of Lemma 4.1

Proof. We are interested in the total derivative of the pass probability with respect to α for any agent with effectiveness $\mu_0 \geq \mu_b$. Note that for this lemma, we consider the agent to always be participating. This total derivative can be expanded using the multi-variable chain rule as follows:

$$\frac{d}{d\alpha}\text{Pass}(\alpha, \mu_0, n_{\mu_0}(\alpha)) = \underbrace{\frac{\partial \text{Pass}}{\partial \alpha}}_{\text{Direct effect}} + \underbrace{\frac{\partial \text{Pass}}{\partial n} \cdot \frac{dn_{\mu_0}(\alpha)}{d\alpha}}_{\text{indirect effect}}. \tag{9}$$

To simplify notation, denote $w_\alpha = \Phi^{-1}(1-\alpha)$, and let

$$\xi(n, \alpha) = \frac{\Phi^{-1}(1-\alpha)\sigma_b - \sqrt{n}(\mu_0 - \mu_b)}{\sigma_0} = \frac{\sigma_b}{\sigma_0}w_\alpha - \frac{\sqrt{n}(\mu_0 - \mu_b)}{\sigma_0}$$

Then we have: $\text{Pass}(\alpha, \mu_0, n_{\mu_0}(\alpha)) = 1 - \Phi(\xi(n_{\mu_0}(\alpha), \alpha))$. We now separate the analysis based on the direct and indirect influence.

Direct effect: The partial derivative with respect to α is:

$$\frac{\partial \text{Pass}}{\partial \alpha} = -\frac{d}{d\alpha}\Phi(\xi(n, a)) = -\phi(\xi) \cdot \frac{\partial \xi}{\partial \alpha} = -\phi(\xi)\frac{\sigma_b}{\sigma_0}\frac{dw_\alpha}{d\alpha} = \frac{\sigma_b}{\sigma_0}\frac{\phi(\xi)}{\phi(z_\alpha)} \geq 0$$

since a higher α lowers the critical threshold w_α and makes the test easier to pass.

Indirect effect: The indirect effect depends on the optimal number of samples the agent uses for a given α. That is, while having more samples increases the chance of passing ($\frac{\partial \text{Pass}}{\partial n} > 0$), the agent might reduce this effort when α increases (since the test becomes easier). The product of these two terms is therefore unclear. Formally,

$$\frac{\partial \text{Pass}}{\partial n} = -\frac{d}{dn}\Phi(\xi(n,\alpha)) = -\phi(\xi)\frac{\partial \xi}{\partial n} = \phi(\xi)\cdot\frac{\mu_0-\mu_b}{2\sigma_0\sqrt{n}} > 0$$

To compute $\frac{dn_{\mu_0}(\alpha)}{d\alpha}$, let $F(n,\alpha)$ be the first-order condition of the agent's utility function: $u(\alpha,\mu_0,n) = R\text{Pass}(\alpha,\mu_0,n) - cn - c_0$. Since we know $n_{\mu_0}(\alpha)$ is the optimal number of samples, we claim that:

$$F(\alpha, n_{\mu_0}(\alpha)) = R\cdot\frac{\partial \text{Pass}}{\partial n}\bigg|_{n=n_{\mu_0}(\alpha)} - c = 0. \tag{10}$$

This holds because, as we show immediately below, the second derivative of the Pass$(\cdot)$ function is always strictly negative with respect to n, meaning the utility function is always strictly concave. This is formalized below:

Lemma C.1. *The second derivative of Pass with respect to n is always negative – $\frac{\partial^2 Pass}{\partial n^2} < 0$.*

Proof. To see this,

$$\frac{\partial^2 \text{Pass}}{\partial n^2} = \frac{d}{dn}\left(\phi(\xi)\cdot\frac{(\mu_0-\mu_b)}{2\sigma_0\sqrt{n}}\right) = \phi(\xi)\cdot\left(-\frac{\xi(\mu_0-\mu_b)^2}{4\sigma_0^2 n} - \frac{(\mu_0-\mu_b)}{4\sigma_0 n^{3/2}}\right).$$

Since $\phi(\xi) > 0$, $\mu_0-\mu_b > 0$, $\sigma_0 > 0$, and $n > 0$, and since ξ can be either sign, the two terms in the parentheses are both negative (even if $\xi < 0$, the negative sign in front ensures negativity). Thus, the entire expression is strictly negative: $\frac{\partial^2 \text{Pass}}{\partial n^2} < 0$. Intuitively this means that every new sample increases the chance of passing, but each one helps less than the last. □

The result above also means that the first derivative of $F(\alpha,n)$ is always non-zero at $n_{\mu_0}(\alpha)$. It is also evident that $F(\alpha,n)$ is continuously differentiable. Thus, we can apply the Implicit Function Theorem. Differentiating both sides of $F(\alpha, n_{\mu_0}(\alpha)) = 0$ with respect to α yields:

$$\frac{\partial F}{\partial n}(\alpha, n_{\mu_0}(\alpha))\cdot\frac{dn_{\mu_0}(\alpha)}{d\alpha} + \frac{\partial F}{\partial \alpha}(\alpha, n_{\mu_0}(\alpha)) = 0 \implies \frac{dn_{\mu_0}(\alpha)}{d\alpha} = -\frac{\frac{\partial F}{\partial \alpha}}{\frac{\partial F}{\partial n}} \tag{11}$$

Next we compute $\frac{\partial F}{\partial \alpha}$ and $\frac{\partial F}{\partial n}$ accordingly. Observe that:

$$\frac{\partial F}{\partial \alpha} = R_0\cdot\frac{\partial^2 \text{Pass}}{\partial n\,\partial\alpha}, \quad \text{where} \quad \frac{\partial^2 \text{Pass}}{\partial n\,\partial \alpha} = \frac{\mu_0-\mu_b}{2\sigma_0\sqrt{n}}\cdot\frac{\partial\phi(\xi)}{\partial\alpha}.$$

Since:

$$\frac{\partial\phi(\xi)}{\partial\alpha} = \frac{d\phi(\xi)}{d\xi}\cdot\frac{\partial\xi}{\partial\alpha} = -\xi\phi(\xi)\cdot\frac{\partial\xi}{\partial\alpha} = \xi\phi(\xi)\cdot\frac{\sigma_b}{\sigma_0}\frac{1}{\phi(w_\alpha)},$$

we have

$$\frac{\partial F}{\partial\alpha} = R_0\cdot\frac{\mu_0-\mu_b}{2\sigma_0\sqrt{n}}\cdot\xi\phi(\xi)\cdot\frac{1}{\phi(z_\alpha)}\frac{\sigma_b}{\sigma_0}.$$

We next turn to computing the derivative of F with respect to n. Observe the following (where we plug in $\frac{\partial^2 \text{Pass}}{\partial n^2}$ from above):

$$\frac{\partial F}{\partial n} = R_0\cdot\frac{\partial^2 \text{Pass}}{\partial n^2} = R_0\cdot\left(\frac{\xi\phi(\xi)(\mu_0-\mu_b)^2}{4\sigma_0^2 n} - \frac{\phi(\xi)(\mu_0-\mu_b)}{4\sigma_0 n^{3/2}}\right).$$

Now plugging them back into the expression computed by the Implicit function theorem, we have:

$$\frac{dn_{\mu_0}(\alpha)}{d\alpha} = -\frac{\frac{\partial F}{\partial\alpha}}{\frac{\partial F}{\partial n}}\bigg|_{n=n_{\mu_0}(\alpha)} = -\frac{\frac{(\mu_0-\mu_b)\,\xi\,\phi(\xi)}{2\sigma_0\sqrt{n_{\mu_0}(\alpha)}\,\phi(w_\alpha)}\frac{\sigma_b}{\sigma_0}}{\frac{\xi\phi(\xi)(\mu_0-\mu_b)^2}{4\sigma_0^2 n_{\mu_0}(\alpha)} - \frac{\phi(\xi)(\mu_0-\mu_b)}{4\sigma_0 n_{\mu_0}(\alpha)^{3/2}}} = -\frac{\frac{2\xi\sqrt{n_{\mu_0}(\alpha)}\sigma_b}{\phi(w_\alpha)}}{\xi(\mu_0-\mu_b) - \frac{\sigma_0}{\sqrt{n_{\mu_0}(\alpha)}}}.$$

Total derivative. Having computed the direct and indirect effects, we can directly compute the total derivative. We have:

$$\frac{d}{d\alpha}\text{Pass}(\alpha,\mu_0,n_{\mu_0}(\alpha)) = \underbrace{\frac{\partial \text{Pass}}{\partial \alpha}}_{\text{Direct effect}} + \underbrace{\frac{\partial \text{Pass}}{\partial n}\Big|_{n=n_{\mu_0}(\alpha)} \cdot \frac{dn_{\mu_0}(\alpha)}{d\alpha}}_{\text{indirect effect}}. \tag{12}$$

$$= \frac{\sigma_b}{\sigma_0}\frac{\phi(\xi)}{\phi(z_\alpha)} + \phi(\xi)\cdot\frac{\mu_0-\mu_b}{2\sigma_0\sqrt{n_{\mu_0}(\alpha)}}\cdot\frac{dn_{\mu_0}(\alpha)}{d\alpha} \tag{13}$$

$$= \frac{\sigma_b}{\sigma_0}\frac{\phi(\xi)}{\phi(z_\alpha)}\left(1-\frac{(\mu_0-\mu_b)\,\xi}{\left(\xi(\mu_0-\mu_b)-\frac{\sigma_0}{\sqrt{n_{\mu_0}(\alpha)}}\right)}\right). \tag{14}$$

The sign of $\left(1-\frac{(\mu_0-\mu_b)\,\xi}{\left(\xi(\mu_0-\mu_b)-\frac{\sigma_0}{\sqrt{n_{\mu_0}(\alpha)}}\right)}\right)$ decides the monotonicity of the passing probability. As we show below, when agent chooses the optimal number of samples $n = n_{\mu_0}(\alpha)$, it will never be the case that $\xi(\mu_0-\mu_b)-\frac{\sigma_0}{\sqrt{n(\alpha)}} > 0$ (Lemma C.2), which implies that $\left(1-\frac{(\mu_0-\mu_b)\,\xi}{\left(\xi(\mu_0-\mu_b)-\frac{\sigma_0}{\sqrt{n_{\mu_0}(\alpha)}}\right)}\right) > 0$ always holds at the optimal $n_{\mu_0}(\alpha)$, thus the passing probability at the optimal number of samples $n_{\mu_0}(\alpha)$ is always monotonically increasing. We finish the proof by proving Lemma C.2:

Lemma C.2. *When $n_{\mu_0}(\alpha)$ is the optimal sample size chosen by the agent to maximize utility, then $\xi(\mu_0-\mu_b)-\dfrac{\sigma_0}{\sqrt{n(\alpha)}} > 0$ never holds.*

Proof. Recall the second derivative of the utility w.r.t n is:

$$\begin{aligned}\frac{\partial^2 \text{Utility}(n,\mu_0;\mu_b,\alpha)}{\partial n^2} &= \frac{\partial^2 \text{Pass}}{\partial n^2}\\ &= \frac{\xi\phi(\xi)(\mu_0-\mu_b)^2}{4\sigma_0^2 n} - \frac{\phi(\xi)(\mu_0-\mu_b)}{4\sigma_0 n^{3/2}}\\ &= \frac{\phi(\xi)(\mu_0-\mu_b)}{4\sigma_0^2 n}(\xi(\mu_0-\mu_b)-\frac{\sigma_0}{\sqrt{n}})\end{aligned}$$

At the optimal number of sample $n = n_{\mu_0}(\alpha)$, the second derivative of the utility must be negative, otherwise it implies that increasing the number of samples will increase the utility, which leads to a contradiction to the fact that $n = n_{\mu_0}(\alpha)$ is an optimal number of samples, thus we have the following always holds at $n = n_{\mu_0}(\alpha)$: $\xi(\mu_0-\mu_b)-\frac{\sigma_0}{\sqrt{n_{\mu_0}(\alpha)}} \le 0$. □

□

Proof of Theorem 4.2

Proof. Recall that the false negative rate in our setting consists of two components, conditioned on whether agents participate. Observing that when an agent does not participate – i.e. $\mu_0 < \mu_\tau(\alpha)$ – their probability of passing the statistical test is 0, we can simplify the overall FN rate as follows:

$$\begin{aligned}\text{FN}(\alpha,\mathcal{I}) &= \mathop{\mathbb{E}}_{\mu_0\sim q}[\text{Fail}(\alpha,\mu_0,n_{\mu_0}(\alpha))|\mu_0\ge\mu_b]\\ &= \underbrace{\mathop{\mathbb{E}}_{\mu_0\sim q}[\text{Fail}(\alpha,\mu_0,n_{\mu_0}(\alpha))|\mu_0\ge\mu_b,\mu_0\ge\mu_\tau(\alpha)]P(\mu_0\ge\mu_\tau(\alpha)|\mu_0\ge\mu_b)}_{\text{FN}_{\text{particip}}} + \underbrace{P[\mu_0\le\mu_\tau(\alpha)|\mu_0\ge\mu_b]}_{\text{FN}_{\text{abstain}}}\end{aligned}$$

Consider first any $\alpha_1 \leq \widehat{\alpha}$. We can explicitly express each component of the false negative as follows, since we are guaranteed that $\mu_\tau(\alpha_1) \geq \mu_b$:

$$\text{FN}_{\text{particip}}(\alpha_1, \mathcal{I}) = \frac{1 - Q(\mu_\tau(\alpha_1))}{1 - Q(\mu_b)} \int_{\mu_\tau(\alpha_1)}^{1} \text{Fail}(\alpha_1, \mu_0, n_{\mu_0}(\alpha)) \frac{q(\mu_0)}{1 - Q(\mu_\tau(\alpha_1))} d\mu_0 \tag{15}$$

$$\text{FN}_{\text{abstain}}(\alpha_1, \mathcal{I}) = \frac{Q(\mu_\tau(\alpha_1)) - Q(\mu_b)}{1 - Q(\mu_b)} \tag{16}$$

Let us focus on $\text{FN}_{\text{particip}}$. To simplify this, the following result is helpful. For a positive function $g(x)$ and a function $f(x)$ with minimum and maximum values f_{min} and f_{max} over an interval $[a, b]$, the following holds:

$$f_{min} \int_a^b g(x) \leq \int_a^b g(x) f(x) dx \leq f_{max} \int_a^b g(x) \implies f_{min} \leq \frac{\int_a^b g(x) f(x) dx}{\int_a^b g(x)} \leq f_{max}$$

Since the expression in the middle is between f_{min} and and f_{max}, by the intermediate value theorem, there exists an $c \in [a, b]$ such that: $f(c) = \frac{\int_a^b g(x) f(x) dx}{\int_a^b g(x)}$ which means: there exists a c such that $f(c) \int_a^b g(x) = \int_a^b f(x) g(x)$. This can be interpreted as an integral mean value theorem. Using this, there exists a value $\mu_c^1 \in [\mu_\tau(\alpha_1), 1]$ such that the $\text{FN}_{particip}(\alpha_1)$ can be expressed as follows:

$$\begin{aligned}
\text{FN}_{\text{particip}}(\alpha_1, \mathcal{I}) &= \frac{1 - Q(\mu_\tau(\alpha_1))}{1 - Q(\mu_b)} \text{Fail}(\alpha_1, \mu_c^1, n_{\mu_c^1}(\alpha)) \int_{\mu_\tau(\alpha_1)}^{1} \frac{q(\mu_0)}{1 - Q(\mu_\tau(\alpha_1))} d\mu_0 \\
&= \frac{1}{1 - Q(\mu_b)} \text{Fail}(\alpha_1, \mu_c^1, n_{\mu_c^1}(\alpha)) \int_{\mu_\tau(\alpha_1)}^{1} \frac{q(\mu_0)}{1 - Q(\mu_\tau(\alpha_1))} d\mu_0 \\
&= \frac{1 - Q(\mu_\tau(\alpha_1))}{1 - Q(\mu_b)} \text{Fail}(\alpha_1, \mu_c^1, n_{\mu_c^1}(\alpha))
\end{aligned}$$

Therefore, we can write:

$$\text{for some } \mu_c^1 \in [\mu_\tau(\alpha_1), 1]: \quad \text{FN}(\alpha_1, \mathcal{I}) = \frac{1 - Q(\mu_\tau(\alpha_1))}{1 - Q(\mu_b)} \text{Fail}(\alpha_1, \mu_c^1, n_{\mu_c^1}(\alpha)) + \frac{Q(\mu_\tau(\alpha_1)) - Q(\mu_b)}{1 - Q(\mu_b)}$$

Now consider an α_2 such that $\alpha_1 < \alpha_2 \leq \alpha^*$. Note that due to Appendix B, $\mu_\tau(\alpha_2) \leq \mu_\tau(\alpha_1)$. We can thus write the $\text{FN}_{\text{particip}}$ of this instance as follows:

$$\begin{aligned}
&\text{FN}_{\text{particip}}(\alpha_2, \mathcal{I}) = \frac{1 - Q(\mu_\tau(\alpha_2))}{1 - Q(\mu_b)} \int_{\mu_\tau(\alpha_2)}^{1} \text{Fail}(\alpha_2, \mu_0, n_{\mu_0}(\alpha)) \frac{q(\mu_0)}{1 - Q(\mu_\tau(\alpha_2))} d\mu_0 \\
&= \frac{1}{1 - Q(\mu_b)} \left[\int_{\mu_\tau(\alpha_2)}^{\mu_\tau(\alpha_1)} \text{Fail}(\alpha_2, \mu_0, n_{\mu_0}(\alpha_2)) q(\mu_0) d\mu_0 + \int_{\mu_\tau(\alpha_1)}^{1} \text{Fail}(\alpha_2, \mu_0, n_{\mu_0}(\alpha_2)) q(\mu_0) d\mu_0 \right] \\
&= \frac{1}{1 - Q(\mu_b)} \left[\text{Fail}(\alpha_2, \mu_c^{2,1}, n_{\mu_c^{2,1}}(\alpha)) \int_{\mu_\tau(\alpha_2)}^{\mu_\tau(\alpha_1)} q(\mu_0) d\mu_0 + \text{Fail}(\alpha_2, \mu_c^{2,2}, n_{\mu_c^{2,2}}(\alpha)) \int_{\mu_\tau(\alpha_1)}^{1} q(\mu_0) d\mu_0 \right] \\
&= \text{Fail}(\alpha_2, \mu_c^{2,1}, n_{\mu_c^{2,1}}(\alpha)) \frac{Q(\mu_\tau(\alpha_1)) - Q(\mu_\tau(\alpha_2))}{1 - Q(\mu_b)} + \text{Fail}(\alpha_2, \mu_c^{2,2}, n_{\mu_c^{2,2}}(\alpha)) \frac{1 - Q(\mu_\tau(\alpha_1))}{1 - Q(\mu_b)}
\end{aligned}$$

where in the third transition, we use the integral mean value theorem as in the previous α_1 case. Note that for the integral between $[\mu_\tau(\alpha_1), 1]$ in the α_2 setting, we know that for every $\mu_0 \in [\mu_\tau(\alpha_1), 1]$, $\text{Fail}(\alpha_2, \mu) < \text{Fail}(\alpha_1, \mu)$ since from Lemma 4.1, we know that the pass probability increases as alpha increases and this set of agents participated under both α_1 and α_2. This immediately means that $\text{Fail}(\alpha_2, \mu_c^{2,2}, n_{\mu_c^{2,2}}(\alpha)) \leq \text{Fail}(\alpha_2, \mu_c^1, n_{\mu_c^{2,1}}(\alpha))$. In the $[\mu_\tau(\alpha_2), \mu_\tau(\alpha_1)]$, the failure probability is at most 1. Hence, $\text{Fail}(\alpha_2, \mu_c^{2,1}, n_{\mu_c^{2,1}}(\alpha)) \leq 1$. Thus, we can upper bound the FN participation loss at α_2 as follows:

$$\text{FN}_{\text{particip}}(\alpha_2) \leq \frac{Q(\mu_\tau(\alpha_1)) - Q(\mu_\tau(\alpha_2))}{1 - Q(\mu_b)} + \text{Fail}(\alpha_1, \mu_c^1, n_{\mu_c^1}(\alpha)) \frac{1 - Q(\mu_\tau(\alpha_1))}{1 - Q(\mu_b)} \tag{17}$$

991 We now express the abstain loss for α_2 as follows:

$$\mathrm{FN}_{\mathrm{abstain}}(\alpha_2) = \frac{Q(\mu_\tau(\alpha_2)) - Q(\mu_b)}{1 - Q(\mu_b)} = \frac{Q(\mu_\tau(\alpha_1)) - Q(\mu_b)}{1 - Q(\mu_b)} - \frac{Q(\mu_\tau(\alpha_1)) - Q(\mu_\tau(\alpha_2))}{1 - Q(\mu_b)}$$

992 Then we have that:

$$\mathrm{FN}(\alpha_2, \mathcal{I}) = \mathrm{FN}_{\mathrm{particip}} + \mathrm{FN}_{\mathrm{abstain}} \tag{18}$$

$$\leq \mathrm{Fail}(\alpha_1, \mu_c^1, n_{\mu_c^1}(\alpha)) \frac{1 - Q(\mu_\tau(\alpha_1))}{1 - Q(\mu_b)} + \frac{Q(\mu_\tau(\alpha_1)) - Q(\mu_b)}{1 - Q(\mu_b)} = \mathrm{FN}(\alpha_1, \mathcal{I}) \tag{19}$$

993 Lastly, for any $\alpha \geq \widehat{\alpha}$, we know from Proposition 4.2 that $\mathrm{FN}_{\mathrm{particip}}(\alpha, I) = 0$ and from Proposi-
994 tion 4.1 that $\mathrm{FN}_{\mathrm{abstain}}(\alpha, \mathcal{I}) = 0$. □

995 D Plot for Vaccine Drugs

Figure 4: $\widehat{\alpha}$ vs Revenue - Vaccines