Incentive Design in Sequential Statistical Protocols

Anonymous Authors[1]

Abstract

Many statistical protocols are deployed inside incentive systems: firms decide whether to run and submit trials, researchers decide which projects to pursue, and regulators or conferences commit to approval rules before observing outcomes. We study a sequential principal-agent model of hypothesis testing in which a principal commits to a history-dependent testing threshold and an agent chooses costly effort and submission. The principal observes only public testing outcomes, not the latent quality that determines social value. Even in this no-feedback environment, dynamic testing rules can create continuation incentives that encourage effort and selective submission. We illustrate this effect in a two-period construction and analyze timeout policies in an infinite-horizon discounted setting, showing how future access to testing can substitute for direct observation of the agent's private information.

1. Introduction

The humble hypothesis test is well-designed for detecting effects through noise, using data generated from scientific experiments, so that its results are influenced solely by the laws of nature. Researchers posit a hypothesis, but reject it whenever the evidence against it is strong, operationalized by comparing a p-value to a fixed threshold. Upon this intellectual framework, hypothesis tests have played a core role in scientific epistemology and evidence accumulation for over a century.

But hypothesis tests have also been deployed in settings whose results are not determined *solely* by the laws of nature. Instead, they are mediated by strategic behavior, by agents who stand to gain from the results of the test. In this light, how might we interpret the p-values of, for example, clinical trials for drugs? It is true that their distribution depends on whether or not the drug works. But if the drugmaker did not believe that the drug would be profitable, then it would not have been submitted in the first place, which should be taken into account. However, the drugmaker's belief in profitability is private, influenced by internal investigations and effort.

Something similar is true for scientific publishing, for example in laboratory sciences where $p < 0.05$ is often a bright line for paper acceptance, and computer science conference culture, where noisy scores generated by imperfect reviewers must, essentially, exceed some threshold. In each case, again the scientist has private information regarding the quality of the project influenced by private effort. In these settings, the behavior of the humble hypothesis test must be understood with respect to the game between strategic agents in which it is embedded.

An important feature of such games is *no-feedback*. In hypothesis testing, when testing a null hypothesis $H_0 : \theta < 0$, the value θ is never ultimately observed. When embedded in a game between the FDA and the drugmaker, θ might represent some drug quality indicator that is not observed during the strategic interaction. Nevertheless, when the agent can affect the distribution of θ through costly effort, dynamic testing protocols can create incentives for the agent to choose higher-quality projects.

Our contribution is to model this interaction as a sequential statistical protocol rather than a one-shot test. The principal commits to a rule for updating future testing thresholds after public approval and denial events. These future thresholds act as continuation incentives: an agent may be willing to incur effort and withhold low-quality submissions today in order to preserve future access to favorable testing conditions.

1.1. Prior work

This paper builds on a growing literature on statistical decision rules with strategic participants. Tetenov (2016) studies regulatory approval as an economic theory of statistical testing, showing how a simple threshold can screen privately informed proponents in a one-shot clinical-trial setting. Bates et al. (2024) formalize principal-agent hypothesis testing, emphasizing that the statistical protocol changes which agents choose to enter. Subsequent work

[1]Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: Anonymous Author <anon.email@domain.com>.

Preliminary work. Under review by the International Conference on Machine Learning (ICML). Do not distribute.

sharpens this connection between testing and incentives: Shi et al. (2024) derive false-discovery guarantees with risk-sensitive agents, Shi et al. (2025) design menus that elicit heterogeneous agent types, and Chen et al. (2025) characterize optimal thresholds when strategic submission affects observed evidence. Our model differs by making the testing protocol sequential and by giving the agent a costly effort decision that changes the distribution of project quality.

The work also relates to statistical decision theory for trials and evidence thresholds. Manski & Tetenov (2021) and Peden & Sprenger (2021) discuss limitations of conventional significance testing as a decision rule, especially when the relevant objective is welfare rather than error control alone. Kasy & Spiess (2022) cast pre-analysis plans as commitment devices in a mechanism-design problem with strategic analysts. We share the view that statistical procedures should be evaluated as incentive-constrained decision rules, but focus on repeated hypothesis tests where the principal can use future access to tests as an incentive instrument.

Finally, our dynamic formulation connects to sequential testing, reputation, and Stackelberg games. Online multiple-testing methods such as α-investing (Foster & Stine, 2008) update future testing levels based on past discoveries, but their objective is error control rather than strategic effort provision. Dynamic moral-hazard and reputation models study how future payoffs discipline hidden actions (Mailath & Samuelson, 2001; Gerardi & Maestri, 2012; Saeedi & Shourideh, 2023); here, the disciplining device is the continuation testing threshold. Computationally, the model is a stochastic Stackelberg game, related to work on discounted stochastic Stackelberg equilibria (Vorobeychik & Singh, 2012; Goktas et al., 2022; 2023), with the special structure that the state is a statistical threshold and the transition is induced by public test outcomes.

2. Online principal-agent hypothesis testing

Problem parameters are $C_w, C_s, \omega_1, \omega_0, R > 0$, a success probability $s \in [0,1]$, two priors $\mathbb{Q}^+$ and $\mathbb{Q}^-$ on $\theta \in \{0,1\}$ with $\mathbb{Q}^+ \succeq \mathbb{Q}^-$, and two p-value densities $f_1(p) = f(p \mid \theta = 1)$ and $f_0(p) = f(p \mid \theta = 0) = 1$ on $p \in [0,1]$. We assume smaller p-values are more likely under the alternative, so $F_1(\alpha) \geq F_0(\alpha)$ for the relevant thresholds.

Summary: A regulator (the principal) declares α_t. Then a drug company (the agent) either works or shirks, and depending on the outcome, updates their confidence $\mathbb{Q}_t$ in $\theta_t \in \{0,1\}$. Subsequently, they decide whether or not to pay a cost to submit their project to a trial, and observe a p-value p_t. A regulator then checks the threshold and issues the decision $\delta \in \{\textsf{approves}, \textsf{denies}\}$.

This generalizes principal-agent hypothesis testing in some ways (Shi et al., 2024), as it adds the sequential and effort components, but unlike that work assumes that the pair $\mathbb{Q}^+$, $\mathbb{Q}^-$ is common knowledge.

2.1. Utilities

Suppose the two players share a discount factor $\gamma \in [0,1]$. The agent's payoff aggregates rewards from approvals net of the costs of effort and submission:

$$U = \sum_{t=1}^{T} \gamma^t \Big(R\,\mathbf{1}\{\delta_t = \textsf{approves}\} - C_w\,\mathbf{1}\{e_t = \textsf{work}\} - C_s\,\mathbf{1}\{a_t = \textsf{opt in}\}\Big).$$

The agent therefore values approval itself, while effort and submission are privately costly.

The principal's objective is to approve high-quality projects and avoid approving low-quality ones. The principal's utility is therefore given by

$$V = \sum_{t=1}^{T} \gamma^t \Big(\omega_1\,\mathbf{1}\{\delta_t{=}\textsf{approves},\ a_t{=}\textsf{opt in},\ \theta_t{=}1\} - \omega_0\,\mathbf{1}\{\delta_t{=}\textsf{approves},\ a_t{=}\textsf{opt in},\ \theta_t{=}0\}\Big).$$

Importantly, the principal's payoff depends on the true state θ_t, even though this state is never observed during the interaction.

The game is Stackelberg: the principal commits ex ante to the approval protocol $\boldsymbol{\alpha}$, and the agent best responds to this commitment. As a result, both the agent's and the principal's expected utilities can be expressed as functions of $\boldsymbol{\alpha}$ alone. Although the true state θ_t may become observable ex post, such information arrives too late to affect incentives or decisions within the game and therefore plays no role in the analysis.

3. A simple illustration in $T = 2$

In what follows, we assume a special case of the above. Suppose $C_s = 0$, and let $\Pr_{\theta \sim \mathbb{Q}^+}(\theta = 1) = 1$ and $\Pr_{\theta \sim \mathbb{Q}^-}(\theta = 0) = 1$, so that the agent *knows* what θ they receive.

The FDA operates by taking $\alpha_1 = \cdots = \alpha_T$, i.e. they set a constant α. The principal can choose $\alpha_1 = \alpha^*$ to optimize their own utility, taking the form

$$\alpha^* = \underset{C_w/R \leq s[F_1(\alpha) - F_0(\alpha)]}{\arg\max} \Big\{ \omega_1 s F_1(\alpha) - \omega_0 (1-s) F_0(\alpha) \Big\}. \tag{1}$$

Call the resulting policy π_{const}. Note that the constraint set, $\{\alpha \in [0,1] : C_w/R \leq s[F_1(\alpha) - F_0(\alpha)]\}$, is precisely the incentive compatibility constraint for the agent to exert effort.

Principal-agent multiple testing

The principal publishes a protocol $\boldsymbol{\alpha} = \alpha_1(\cdot), \alpha_2(\cdot), \ldots$, which are functions of the common knowledge at time t, generically denoted as $\mathcal{H}_t$.
Then, the following stage game is played at times $t = 1, \ldots, T$:

- The agent decides $e_t \in \{\textsf{work}, \textsf{shirk}\}$.
- If work, they pay cost C_w, and their effort succeeds with probability s. They privately observe whether or not their effort succeeds.
  - If their effort succeeds, then $\mathbb{Q}_t = \mathbb{Q}^+$.
  - Otherwise, $\mathbb{Q}_t = \mathbb{Q}^-$.
- If shirk then $\mathbb{Q}_t = \mathbb{Q}^-$.
- The agent then decides whether $a_t \in \{\textsf{opt in}, \textsf{opt out}\}$.
- If opt in, they pay cost C_s, and:
  - Nature draws $\theta_t \sim \mathbb{Q}_t$, which is not observed.
  - The agent draws a statistic $p_t \sim f_{\theta_t}$, which is common knowledge.
  - The principal decides $\delta_t = \textsf{approves}$ if $p_t \leq \alpha_t(\mathcal{H}_t)$, and else denies.
  - If approves, then the agent receives reward R, and at the end of the game, the principal receives ω_1 if $\theta_t = 1$ and $-\omega_0$ if $\theta_t = 0$. (They do not observe these rewards, or θ_t, during the game.)

Figure 1. The sequential principal-agent hypothesis testing protocol.

The constant policy may be optimal—not too surprising, given that there is a vast space of non-constant policies—but we only need a simple policy to demonstrate this. Namely, consider the following policy: for odd t, $\alpha_t = \alpha^*$, and for even t, $\alpha_t = 0$ if $\delta_t = \textsf{denies}$, and otherwise $\alpha_t = \alpha^*$. Call this policy π_{alt}.

Theorem 3.1. *Suppose that the pair f_1 and $f_0 = 1$ satisfy the monotone-likelihood ratio condition, as well as the technical condition*

$$\lim_{x \to 0^+} \frac{F_1(x)}{x f_1(x)} < 2.$$

Then for any such pair, there exist parameters $C_w, R, \omega_1, \omega_0 > 0$ such that the principal's expected utility under π_{alt} is strictly larger than under π_{const} when the agent best responds. More precisely, let $\rho := C_w/R$, and $\eta = \omega_0/\omega_1$. Then there exists a value $\underline{\eta} > 0$ and a function $\bar{\rho}(\eta) > 0$ such that the strict improvement holds whenever $\eta \in [\underline{\eta}, \infty)$ and $\rho \in [0, \bar{\rho}(\eta))$.

Proof. In this proof, let $C_w = C$. Now suppose the claim is true for $T = 1$ and $T = 2$. Then it is true for any $T = K$, because if K is even, then the $T = K$ game consists of $K/2$ copies of the $T = 2$ game, and if K is odd, then it consists of $\lfloor K/2 \rfloor$ copies of the $T = 2$ game and one copy of the $T = 1$ game.

Additionally, the proof for $T = 1$ is trivial (the two policies give the same result). Let us now prove it for $T = 2$.

Given $\boldsymbol{\alpha}$ and problem parameters, let us say the agent does action A1 if their best response in the first round can always be $e_1 = \textsf{work}$ and $a_1 = \textsf{opt in}$; they are action A2 if instead $e_1 = \textsf{work}$ and

$$a_1 = \begin{cases} \textsf{opt in} & \theta_1 = 1 \\ \textsf{opt out} & \theta_1 = 0 \end{cases}.$$

and they take action B if $e_1 = \textsf{shirk}$. Due to incentive compatibility, under π_{const} using α^* the agent must take action A1 or A2 (in fact, A1). It follows that under π_{alt} the agent must also take action A1 or A2: action B is worse in the first round by the choice of α^*, and worse in the second round because shirking reduces the chance of reaching the favorable continuation state. Hence, without loss of generality, we ignore agents who take action B.

Let the principal's expected utility under some policy π be $\mathcal{V}(\pi)$, and let the expected utility for an agent who does action Q which best responds in the second round be $\mathcal{V}(\pi; \text{Q})$. Let

$$\begin{aligned} \mathcal{V}_1(\alpha^*) &= \omega_1 s \Pr(p_1 < \alpha^*, a_1 = \textsf{opt in}, \theta_1 = 1) \\ &\quad - \omega_0 (1 - s) \Pr(p_1 < \alpha^*, a_1 = \textsf{opt in}, \theta_1 = 0) \\ &= \omega_1 s F_1(\alpha^*) - \omega_0 (1 - s) F_0(\alpha^*) \end{aligned}$$

be the expected principal's utility when the agent does action A1 and the principal uses the constant threshold α^*. With this notation, we have

$$\mathcal{V}(\pi_{\text{const}}, \text{A1}) = 2\mathcal{V}_1(\alpha^*)$$

and

$$\begin{aligned}\mathcal{V}(\pi_{\text{alt}}, \text{A1}) = \mathcal{V}_1(\alpha^*) \\ + (sF_1(\alpha^*) + (1-s)F_0(\alpha^*))\mathcal{V}_1(\alpha^*) \\ < \mathcal{V}(\pi_{\text{const}}, \text{A1}),\end{aligned}$$

so only by incentivizing agents to do action A2 can π_{alt} possibly be better. Let us then derive an inequality which guarantees that

$$\mathcal{V}(\pi_{\text{const}}, \text{A1}) < \mathcal{V}(\pi_{\text{alt}}, \text{A2}).$$

Because

$$\mathcal{V}(\pi_{\text{alt}}, \text{A2}) = \omega_1 sF_1(\alpha^*)\big(1 + sF_1(\alpha^*)\big),$$

(there is probability $sF_1(\alpha^*)$ of an approval in each round), this is equivalent to

$$\begin{aligned}2\big[\omega_1 sF_1(\alpha^*) - \omega_0(1-s)F_0(\alpha^*)\big] \\ < \omega_1 sF_1(\alpha^*)\big(1 + sF_1(\alpha^*)\big),\end{aligned}$$

or, after rearranging,

$$\left(\frac{1-s}{s}\right)\eta > \frac{F_1(\alpha^*)\big(1 - sF_1(\alpha^*)\big)}{2F_0(\alpha^*)}. \tag{2}$$

We will return to (2) later. Next, we will derive the inequality for which an agent's best response is for them to do action A2.

Let an agent of action Q's expected utility under π_{alt} be $\mathcal{U}(\pi_{\text{alt}}; \text{Q})$, and let

$$\mathcal{U}_1(\alpha^*) = R[sF_1(\alpha^*) + (1-s)F_0(\alpha^*)] - C.$$

Action A2 dominates action A1 if

$$\mathcal{U}(\pi_{\text{alt}}; \text{A1}) < \mathcal{U}(\pi_{\text{alt}}; \text{A2})$$

which is equivalent to

$$\begin{aligned}\mathcal{U}_1(\alpha^*) + \Big(sF_1(\alpha^*) + (1-s)F_0(\alpha^*)\Big)\mathcal{U}_1(\alpha^*) \\ < \Big(\mathcal{U}_1(\alpha^*) - RF_0(\alpha^*)(1-s)\Big) \\ + \Big(sF_1(\alpha^*) + (1-s)\Big)\mathcal{U}_1(\alpha^*)\end{aligned}$$

which, after rearranging becomes

$$s[F_1(\alpha^*) - F_0(\alpha^*)] \geq \rho + \frac{(F_0(\alpha^*))^2}{1 - F_0(\alpha^*)}. \tag{3}$$

Now, we must find parameter configurations such that (2) and (3) are simultaneously satisfied, for $\rho := C/R$ and $\eta = \omega_0/\omega_1$.

Let $\alpha^*_{\text{unconst}}(\eta) > 0$ be that which solves $dF_1(\alpha)/dF_0(\alpha) = f_1(\alpha)/f_0(\alpha) = \eta\frac{(1-s)}{s}$. Observe that whenever (3) is satisfied with $\alpha^*_{\text{unconst}}$ in place of α^*, then $\alpha^*_{\text{unconst}} = \alpha^*$, i.e. the constraint in (1) does not bind, and hence it suffices to find parameters such that (3) is satisfied with $\alpha^*_{\text{unconst}}$.

We claim that, for η sufficiently large, we simultaneously satisfy (2), as well as

$$s[F_1(\alpha^*_{\text{unconst}}(\eta)) - F_0(\alpha^*_{\text{unconst}}(\eta))] \geq \frac{(F_0(\alpha^*_{\text{unconst}}(\eta)))^2}{1 - F_0(\alpha^*_{\text{unconst}}(\eta))}. \tag{4}$$

To see why, observe that $\alpha^*_{\text{unconst}}(\eta)$ is decreasing in η, because monotone likelihood ratio implies concavity of $F_1(\alpha)$, and must go to zero as η gets large, because $f_1(\alpha)$ is non-negative and differentiable on $(0,1)$. Then we can take η large to get $\alpha^*_{\text{unconst}}(\eta)$ small, whence (2) is satisfied by the technical assumption and (4) is satisfied because $\alpha^2/(1-\alpha)$ has slope zero at $\alpha = 0$, and the left-hand side has positive slope at $\alpha = 0$ because F_1 stochastically dominates F_0. Let $\underline{\eta}$ be the smallest such η such that (2) and (4) are satisfied.

Then finally for any $\eta \in (\underline{\eta}, \infty)$, we can pick

$$\begin{aligned}\rho = s[F_1(\alpha^*_{\text{unconst}}(\eta)) - F_0(\alpha^*_{\text{unconst}}(\eta))] \\ - \frac{(F_0(\alpha^*_{\text{unconst}}(\eta)))^2}{1 - F_0(\alpha^*_{\text{unconst}}(\eta))}.\end{aligned}$$

Any such pair of ρ, η works. □

3.1. The technical assumption

We have assumed above the following assumption.

Assumption 3.2. The alternative p-value distribution F_1 satisfies

$$\lim_{x \to 0^+} \frac{F_1(x)}{xf_1(x)} < 2.$$

This holds whenever F_1 is differentiable at 0, so that $f_1(0)$ is finite, in which case the left-hand side equals 1. However, in most typical cases in testing, $f_1(0) = \infty$.

Nevertheless, the condition often holds, for example under the z-test when this quantity is equal to 1. The following proposition illustrates a basic model when it may hold.

Proposition 3.3. *Assume that under the null, an underlying test statistic satisfies $T \sim g_0(t)$, and under the alternative $T \sim \exp\{\theta t - A(\theta)\}g_0(t) =: g_\theta(t)$, so that for the p-value testing $H_0 : \theta \leq 0$ against $H_1 : \theta > 0$, we have $p(t) = \bar{G}_0(t) = 1 - G_0(t)$. Then F_1, the distribution of $p(T)$ under the alternative, satisfies*

$$\lim_{x \to 0^+} \frac{F_1(x)}{xf_1(x)} = 1$$

whenever the hazard $h_0(t) := g_0(t)/\bar{G}_0(t)$ is monotone increasing. In particular, it suffices to assume that g_0 is *log-concave, which implies that the hazard rate is monotone increasing.*

Proof. We will show this claim for general $\theta > 0$. Rewrite

$$\frac{F_\theta(p)}{pf_\theta(p)} = \frac{\bar{G}_\theta(t(p))/g_\theta(t(p))}{\bar{G}_0(t(p))/g_0(t(p))} = \mathop{\mathbb{E}}_{T\sim g_0}\left[e^{\theta(T-t(p))} \mid T > t(p)\right],$$

where $t(p) = \bar{G}_0^{-1}(p)$ is monotonic non-decreasing in p, and satisfies $\lim_{p\to 0^+} t(p) = \infty$. It hence suffices to compute the limit

$$\lim_{p\to 0^+} \mathop{\mathbb{E}}_{T\sim g_0}\left[e^{\theta(T-t(p))} \mid T > t(p)\right] = \lim_{t\to\infty} \mathop{\mathbb{E}}_{T\sim g_0}\left[e^{\theta(T-t)} \mid T > t\right].$$

Let U_t be the non-negative random variable with the same distribution as $T - t$, under g_0, conditional on $\{T \geq t\}$; for $u > 0$ it has survival function

$$\Pr(U_t > u) = \frac{\bar{G}_0(t+u)}{\bar{G}_0(t)}. \tag{5}$$

Additionally, the relation $f'(t)\cdot u \geq f(t+u) - f(t)$ holds whenever $f'(t)$ is monotone decreasing. Then, via the relation $\frac{d}{dt}\log \bar{G}_0(t) = -g_0(t)/\bar{G}_0(t) = -h_0(t)$, we get the upper bound

$$\frac{\bar{G}_0(t+u)}{\bar{G}_0(t)} \leq e^{-uh_0(t)}. \tag{6}$$

Now we can compute the limit. By definition of U_t,

$$\lim_{t\to\infty} \mathop{\mathbb{E}}_{T\sim g_0}\left[e^{\theta(T-t)} \mid T > t\right] = \lim_{t\to\infty} \mathbb{E}[e^{\theta U_t}],$$

and so

$$\begin{aligned}\lim_{t\to\infty} \mathbb{E}[e^{\theta U_t}] &= \lim_{t\to\infty} \int_0^\infty \Pr(e^{\theta U_t} > x)\,dx \\ &= 1 + \lim_{t\to\infty} \int_1^\infty \Pr(e^{\theta U_t} > x)\,dx \\ &= 1 + \lim_{t\to\infty} \int_1^\infty \theta e^{\theta u} \Pr(U_t > u)\,du \\ &\leq 1 + \lim_{t\to\infty} \int_1^\infty \theta e^{\theta u} e^{-uh_0(t)}\,du \\ &\leq 1,\end{aligned}$$

where the second line uses $\theta, U_t > 0$, the fourth applies (5) and (6), and in the last line we apply dominated convergence, as the integrand is dominated by θe^{-u} for large enough t. □

4. Discussion

The two-period example and the infinite-horizon timeout calculation make the same point from complementary angles. A fixed threshold can incentivize effort only by making approvals sufficiently likely, but that also exposes the principal to false approvals when the agent submits low-quality projects. A dynamic protocol can separate these incentives: it can make effort valuable through future access while making low-quality submission costly through future denial states.

The timeout family is not meant to exhaust the space of dynamic protocols. Rather, it is a transparent witness that no-feedback sequential testing can differ qualitatively from one-shot principal-agent testing. More general protocols can be computed by backward induction on a discretized threshold grid, using strong Stackelberg tie-breaking to select the principal-preferred action among the agent's best responses. This provides a numerical route for comparing constant thresholds, timeout rules, and richer standing-index policies.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

References

Bates, S., Jordan, M. I., Sklar, M., and Soloff, J. A. Principal-Agent Hypothesis Testing, April 2024.

Chen, Y., Hossain, S., and Chen, Y. Strategic Hypothesis Testing. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, October 2025.

Foster, D. P. and Stine, R. A. α-Investing: A Procedure for Sequential Control of Expected False Discoveries. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 70(2):429–444, April 2008. ISSN 1369-7412. doi: 10.1111/j.1467-9868.2007.00643.x.

Gerardi, D. and Maestri, L. A principal–agent model of sequential testing. *Theoretical Economics*, 7(3):425–463, 2012. ISSN 1555-7561. doi: 10.3982/TE914.

Goktas, D., Zhao, S., and Greenwald, A. Zero-Sum Stochastic Stackelberg Games. In *Advances in Neural Information Processing Systems*, October 2022.

Goktas, D., Prakash, A., and Greenwald, A. Convex-Concave Zero-Sum Markov Stackelberg Games. In *Thirty-Seventh Conference on Neural Information Processing Systems*, November 2023.

Kasy, M. and Spiess, J. Rationalizing Pre-Analysis Plans:Statistical Decisions Subject to Implementability. Economics Series Working Papers 975, University of Oxford, Department of Economics, June 2022.

Mailath, G. J. and Samuelson, L. Who Wants a Good Reputation? *The Review of Economic Studies*, 68(2): 415–441, April 2001. ISSN 0034-6527, 1467-937X. doi: 10.1111/1467-937X.00175.

Manski, C. F. and Tetenov, A. Statistical Decision Properties of Imprecise Trials Assessing Coronavirus Disease 2019 (COVID-19) Drugs. *Value in Health: The Journal of the International Society for Pharmacoeconomics and Outcomes Research*, 24(5):641–647, May 2021. ISSN 1524-4733. doi: 10.1016/j.jval.2020.11.019.

Peden, W. and Sprenger, J. *Statistical Significance Testing in Economics*, pp. 423–432. Routledge, New York, 1 edition, October 2021. ISBN 978-1-315-73979-3. doi: 10.4324/9781315739793-37.

Saeedi, M. and Shourideh, A. Optimal Rating Design under Moral Hazard, July 2023.

Shi, F. C., Bates, S., and Wainwright, M. J. Sharp Results for Hypothesis Testing with Risk-Sensitive Agents, December 2024.

Shi, F. C., Wainwright, M. J., and Bates, S. Instance-Adaptive Hypothesis Tests with Heterogeneous Agents, October 2025.

Tetenov, A. An economic theory of statistical testing. Technical report, The IFS, September 2016.

Vorobeychik, Y. and Singh, S. Computing Stackelberg Equilibria in Discounted Stochastic Games. *Proceedings of the AAAI Conference on Artificial Intelligence*, 26(1): 1478–1484, September 2012. ISSN 2374-3468, 2159-5399. doi: 10.1609/aaai.v26i1.8234.