

Designing Incentives for Academic Research

Nathan Yoder*

University of Wisconsin-Madison

November 18, 2016

[Most recent version here](#)

Abstract

Recent evidence has called into question the reproducibility of published research in the social and biomedical sciences. The scientific community's response has focused on making the research process more transparent and correctly interpreting its results. I focus on a different issue: Even if these reforms are successful, the ways researchers choose to perform their experiments will determine the quantity and accuracy of the positive results they generate. These choices depend on the incentives researchers face.

I consider the problem faced by a research institution which designs these incentives. The institution contracts with a researcher of unknown quality, who can choose a costly experiment to conduct. The verifiable result of this experiment helps a mass of heterogeneous practitioners decide whether to take a risky action. The institution seeks to maximize the practitioners' surplus net of transfers to the researcher. In keeping with current practice, the institution contracts based on the experiment's result instead of its methodology. This removes a degree of freedom from the optimal design problem, but I show that there need not be loss from doing so. The optimal contract has two general characteristics. First, to discourage the production of false positive results, negative results supporting conventional wisdom must be rewarded. Second, the most informative results must be disproportionately rewarded. To arrive at these conclusions, I contribute to the literature by characterizing solutions and comparative statics of Bayesian persuasion problems using differentiability.

* I am grateful to Drew Fudenberg, Marina Halac, Josh Miller, Michael Ostrovsky, Antonio Penta, Dan Quint, Marzena Rostek, Bill Sandholm, Lones Smith, Marek Weretka, and Muhamet Yildiz for their helpful suggestions.

1 Introduction

There is ample concern in the social and biomedical sciences that existing research is not accurate enough to be relied upon by practitioners. Prompted by the low rate of success of large-scale replication studies such as Open Science Collaboration (2015), scholars in these disciplines have begun to consider the possibility that the publication process is producing too many false positive results.

The causes they have proposed for these problems fall into two major categories. One is that referees and editors do not consider how a result was obtained when interpreting it. Instead, this line of criticism goes, they take any p -value of less than .05 as conclusive evidence for an effect's existence. Results of dubious predictive value thus slip into the literature. This critique has been advanced most notably by Ioannidis (2005) and imported to economics by Maniadis et al. (2014).

The other popular explanation is one of moral hazard. Most academic disciplines reward results which overturn conventional wisdom much more than those which confirm it. If methodology is unverifiable, this creates an incentive to tweak the way the data are analyzed until a superficially iconoclastic result is obtained. Simmons et al. (2011) call this utilizing *researcher degrees of freedom*. It is also frequently referred to as *p-hacking*.

These are important criticisms, and the scientific community has already begun to address them. In the United States, the methodology of every clinical trial must now be registered with the federal government before data collection begins.¹ Pre-registration is also now offered by many top psychology journals. The American Statistical Association has issued a formal statement about the proper use of p -values,² and a psychology journal³ has even banned their use in its pages.

But even if institutions succeed in making the research process fully transparent and correctly interpret every result, other issues will linger. Researchers will still choose the informativeness of their experiments in response to the incentives they face. Moreover, they will choose *how* their experiments inform, i.e., the distributions of posterior beliefs they induce. Finally, since researchers' ability levels are largely unobservable, their choices will exhibit adverse selection.

This paper asks what these incentives should look like, and how they should be implemented, to best further the missions of the institutions that support academic research. Specifically, *how can these institutions maximize the surplus society gets from research, net of the cost of rewarding researchers?*

The environment in which I formalize this design problem is as follows. First, a researcher of unknown quality chooses a costly experiment to conduct about the state of the world. Then, a unit mass of practitioners uses the result, which is verifiable, to decide whether to take a risky action. The institution's problem is thus to design a contract for the researcher so as to maximize the practitioners' expected surplus from better decisions minus the contract's expected cost.

¹See 42 U.S.C. §282(j).

²Wasserstein and Lazar (2016).

³*Basic and Applied Social Psychology*.

There are two chief sources of novelty in this setting. One is the choice variable: I allow the researcher to choose *any* experiment which is possible in principle. Naturally, this complicates the analysis. However, the experiment choice problems in this paper each reduce to the Bayesian persuasion problem of Kamenica and Gentzkow (2011). These amount to choosing a probability distribution over posterior beliefs to maximize the expectation of a value function of those beliefs. Unlike real-valued concave optimization problems, their solutions cannot be characterized by standard first-order conditions. Instead, I develop necessary (Theorem 1) and sufficient (Theorems 2 and 5) conditions on the *tangent line* to the value function which are their Bayesian persuasion analogues. These are new to the persuasion literature.⁴ This allows me to exploit differentiability to pin down the optimal experiment as the solution to a system of equations. The implicit function theorem then tells us how the solution changes as the value function becomes more convex (Lemma 12). These results extend to *all* binary-state Bayesian persuasion problems.

The second novel element of my setting is that the choice variable may be difficult to contract on directly. Instead, the current norm in most disciplines is to reward researchers based on the way their results change our beliefs about the state of the world. The revelation principle tells us that the institution can do no better than by contracting on the choice variable. Contracting on results instead constrains the institution in two substantive ways. First, not every payment rule for experiments can be expressed as a payment rule for results. Second, contracting on results creates an ex post participation constraint: if the institution penalizes the researcher for certain results, the researcher can decline to reveal them. However, Theorem 7 shows that these constraints need not bind. Instead, it provides broad conditions under which the institution can achieve the second-best outcome with a results-based contract.

Which results, then, should institutions reward? As mentioned above, negative empirical findings (i.e., those which provide evidence against a novel hypothesis) are hardly rewarded in academia. This might seem intuitive if they are not valuable for decision-making: knowing that an experimental drug does not work is not useful for a doctor who would never think to prescribe it. But if only positive results are rewarded, researchers have an incentive to conduct experiments which produce more false positive results. This makes their positive results less informative — and hence less valuable. Theorem 7 shows that this effect is large enough to make rewarding informative negative results worthwhile, *even when they have no intrinsic value*.

Finally, how should different results be rewarded? Many disciplines — including ours — place outsize importance on publishing in “top N ” journals.⁵ This creates an incentive structure which rewards the strongest positive results⁶ disproportionately. Theorem 7 shows that this makes sense: disproportionately rewarding the most informative results is the best way for institutions to combat adverse selection. However,

⁴In a recent NBER Working Paper, Caplin and Dean (2013) independently formulate tangent line conditions similar to Theorems 1 and 5 in the context of a rational inattention model. (See Lemma 3 in their appendix.)

⁵For instance, $N = 5$ in both economics (*American Economic Review*, *Econometrica*, *Journal of Political Economy*, *Quarterly Journal of Economics*, *Review of Economic Studies*) and medicine (*Annals of Internal Medicine*, *Journal of the American Medical Association*, *New England Journal of Medicine*, *British Medical Journal*, and *The Lancet*).

⁶E.g., those which are informative enough to merit publication in a top N journal.

unlike current practice, this should include disproportionately rewarding highly informative *negative* results!

I motivate my model and results using academic research. However, they apply equally well to any setting in which a principal contracts with a disinterested agent of unknown quality to verifiably conduct an experiment of her choice. Theorem 7 is especially meaningful when it is less difficult to contract on an experiment’s results than on its methodology. Examples include the state hiring prosecutors and police to investigate crime and firms contracting with R&D workers.

When considered at this level of generality, my paper represents a novel contribution to the literature on *principal-expert* problems. In this literature, the closest paper to mine is Osband (1989), which considers a related setting with adverse selection and moral hazard. I innovate upon his environment by explicitly modeling experimentation and allowing experimenters to choose *any* experiment possible in principle, rather than a precision level. More recently, Zermeño (2011) and Carroll (2013) consider moral hazard settings closely related to the adverse selection problem I consider.

Concurrent theoretical work also investigates the impact of incentives on the production of faulty research. Libgober (2015) uses a stylized model with moral hazard to examine the role of transparency requirements, and finds that they may not increase welfare. Di Tillio et al. (2016) show concretely how incentives can lead to hidden researcher actions which influence the outcomes of their experiments.⁷ Outside economics, Smaldino and McElreath (2016) develop an evolutionary model in which labs that produce more positive results are more likely to pass on their methodologies. They show that this quickly results in high false positive rates.

The literature on experimental design following Kamenica and Gentzkow (2011) is large and growing. The same authors consider costly experiments in Gentzkow and Kamenica (2014) and competition between experimenters in Kamenica and Gentzkow (2016). In a binary-action persuasion setting, Kolotilin (2015) gives welfare comparative statics results in the prior distribution on a continuum state space. In a similar setting, Kolotilin et al. (2015) consider the problem of designing an experiment to influence a privately informed receiver. More related to my setting is Bergemann et al. (2016), who consider the problem of selling experiments. Their contracting problem is essentially the reverse of mine: their researcher writes a contract to extract value from an information user with private information, instead of the other way around.

The paper is structured as follows. Section 2 describes the environment in detail. Section 3 solves for the optimal direct revelation contract. The Bayesian persuasion characterization results which I develop to do so lie in Section 3.2. Section 4 shows how to implement outcomes by contracting on experimental results. Section 4.2 characterizes the optimal results-based contract, and Section 5 discusses the main results. All proofs are in the appendix.

⁷Di Tillio et al. (2016) present this as a problem of private information about which of two experiments is biased. However, the friction in their model is due to the fact that the researcher’s choice between the biased and unbiased experiment is hidden.

2 Environment

There is a state of the world $\omega \in \{0, 1\}$; a *researcher*, who is capable of gathering information about ω ; *practitioners*, who make decisions based on their beliefs about ω ; and an *institution*, which contracts with the researcher to benefit the practitioners. Each of these agents places the same prior probability p_0 on the event $\omega = 1$.

The Researcher

The researcher's *type* $\theta \in \Theta = [0, 1]$ describes how skilled she is at investigating the state of the world. Her type is distributed according to the strictly positive density f , and is the researcher's private information.

She may investigate the state of the world ω by choosing a costly experiment $\nu \in X$ to undertake. This experiment produces *hard information*: its realization is verifiable and thus cannot be falsely reported by the researcher. I represent experiments as the *distributions* of posterior beliefs p they induce.⁸ Accordingly, I call an experiment's realized posterior belief its *result*.

All experiments possible in principle are available to the researcher. Thus, X is the set of all distributions of posteriors that can be induced by an experiment. Because Bayesian updating is a martingale process, each $\nu \in X$ must have $E_\nu p = p_0$. Kamenica and Gentzkow (2011) call this property *Bayes-plausibility* and show that it is the *only* restriction on ν implied by the fact that some experiment induces it; hence, $X = \{\nu \in \Delta([0, 1]) : E_\nu p = p_0\}$.

The cost $c(\nu, \theta)$ to a type- θ researcher of undertaking the experiment ν is given by the expected change in a strictly concave measure of uncertainty H which is less concave for higher θ :⁹

$$c(\nu, \theta) = E_\nu [H(p_0, \theta) - H(p, \theta)].$$

I assume that H is four times continuously differentiable (and so $H_{pp}(p, \theta) < 0$ and $H_{pp\theta}(p, \theta) > 0 \forall \theta$), H_θ is bounded, and the marginal cost of certainty is arbitrarily high (and so $\lim_{p \rightarrow 0} H_p(p, \theta) = \lim_{p \rightarrow 1} -H_p(p, \theta) = \infty$).¹⁰ One familiar example of a measure of uncertainty which satisfies these assumptions is *entropy*:

$$H(p, \theta) = -g(\theta)(p \log(p) + (1 - p) \log(p)), \quad g(\theta) > 0, g'(\theta) < 0.$$

I will return to this measure of uncertainty in examples throughout the paper.

Because of Jensen's inequality, the cost function's form ensures that more informative experiments (in the Blackwell sense) are more costly to generate. That is, if ν' is a mean-preserving spread of ν , then

⁸This is without loss — the informational content of an experiment is fully characterized by the distribution of posteriors it induces.

⁹A fixed cost of doing research does not qualitatively change the results; instead, it only affects the participation bound which defines the low type $\underline{\theta}$ (see Lemma 6).

¹⁰In Gentzkow and Kamenica (2014), the authors alternatively assume that c is proportional to the reduction in uncertainty for some agent with prior p_H , which need not be equal to p_0 , so that the cost of an experiment does not depend on the prior of the agent who views it. They then show that we can write c as an expectation with respect to ν of $H(p_H) - \hat{H}(p)$ for appropriate \hat{H} . As long as \hat{H} is itself concave, their functional form assumption is equivalent to the one I make here, with the addition of a fixed cost.

$c(\nu', \theta) \geq c(\nu, \theta)$. Further, the difference in the costs of ν' and ν is smaller for better researchers (those with higher θ): $c(\nu', \theta') - c(\nu, \theta') \geq c(\nu', \theta) - c(\nu, \theta)$ for $\theta' \leq \theta$.

The researcher has von Neumann-Morgenstern preferences, and her Bernoulli utility u is quasilinear in the transfer τ she receives from the institution: $u(\nu, \theta, \tau) = \tau - c(\nu, \theta)$.

Practitioners

After observing the result of the experiment, each of a unit mass of practitioners $i \in [0, 1]$ chooses whether to take a risky action which pays off only in one state of the world $\omega = 1$, or a safe action whose payoff does not depend on the state of the world. Practitioners' outside options are heterogeneous: the safe action's payoff varies across i . When $\omega = 1$, the risky action yields a payoff of ρ ; when $\omega = 0$, it yields a payoff of zero. The safe action pays off $s(i) < \rho$ regardless of ω . Thus, practitioner i takes the risky action if

$$p\rho > s(i) \leftrightarrow p > \frac{s(i)}{\rho}$$

For simplicity, I assume that $s(i) > p_0\rho$ for all i ; that is, no practitioners take the risky action unless they observe information suggesting that $\omega = 1$. In keeping with the language of statistical testing, we can think of the event $\omega = 0$ as the *null hypothesis* and the event $\omega = 1$ as the *alternative hypothesis*. Accordingly, I refer to results $p \leq p_0$ as *negative results* and $p > p_0$ as *positive results*.

For tractability, I assume that s is measurable and that its distribution $\Phi : [\underline{s}, \bar{s}] \rightarrow [0, 1]$ has a continuously differentiable density ϕ . Then practitioner surplus is given by

$$w(p) \equiv \int_{\underline{s}}^{\rho p} (\rho p - s)\phi(s)ds,$$

since the increase in expected utility of practitioner i from viewing the result p is $\rho p - s(i)$ if $s(i) < \rho p$ and zero otherwise. Integration by parts yields

$$w(p) = \int_{\underline{s}}^{\rho p} \Phi(s)ds.$$

Note that w is convex: $w''(p) = \rho^2 p(\rho p)$. Thus, more informative experiments never decrease practitioner welfare.

Institution and Contracting

The institution has von Neumann-Morgenstern preferences with Bernoulli payoffs given by $w(p) - \tau$, practitioner surplus minus transfers τ to the researcher. It works to maximize these by making the researcher a take-it-or-leave-it offer of one of two types of contracts.

A *methods-based contract* consists of a transfer function $T : X \rightarrow \mathbb{R}$. Under this type of contract, when the researcher conducts experiment ν , she is paid $T(\nu)$. The revelation and taxation principles ensure that this class of contracts is equivalent to the class of direct revelation contracts $(\chi : \Theta \rightarrow X, \tau : \Theta \rightarrow \mathbb{R})$.

A *results-based contract* consists of a tariff $\psi : [0, 1] \rightarrow \mathbb{R}_+$. Under this type of contract, when the researcher produces p as the result of some experiment, she is paid $\psi(p)$. Note that the researcher has an ex post participation constraint in this type of contract: while she cannot falsely report results, she need not report a result that would give her a utility penalty.

2.1 Model Discussion

In this paper, I conceptualize an experiment as a distribution of subjective posterior beliefs. Another, perhaps more familiar way to represent an experiment is as a pair of signal distributions $\{\sigma(\cdot|0), \sigma(\cdot|1)\}$ on an outcome space S (which need not be larger than $[0, 1]$). Up to a reparameterization of the outcome space, ν is consistent with a unique pair of signal distributions on $[0, 1]$: From Bayes' rule,

$$\begin{aligned}\sigma(A|1) &= \frac{\int_A p d\nu(p)}{p_0} = \frac{E_\nu[p|p \in A]}{p_0}, \\ \sigma(A|0) &= \frac{\int_A (1-p) d\nu(p)}{1-p_0} = \frac{\nu(A) - E_\nu[p|p \in A]}{1-p_0}\end{aligned}$$

This representation can help us better understand ν by using the language of classical hypothesis testing. When an agent with prior p_0 views the realization of σ , he determines which direction to move his belief as if he were conducting a hypothesis test with critical value p_0 . This is consistent with labeling results $p > p_0$ as positive results and $p \leq p_0$ as negative results. Accordingly, when σ yields a positive result $p > p_0$ even though $\omega = 0$, we can think of p as being a false positive result, or type I error. Likewise, when σ yields a negative result $p \leq p_0$ even though $\omega = 1$, we can think of p as being a false negative result, or type II error. The experiment's type II error rate (i.e., the probability that it gives a negative result when $\omega = 1$) is given by $\beta \equiv \sigma([0, p_0]|1) = 1 - E_\nu[p|p > p_0]/p_0$; its type I error rate is given by $\alpha \equiv \sigma((p_0, 1]|0) = (\nu((p_0, 1]) - E_\nu[p|p > p_0])/(1 - p_0)$.

Then we can write the subjective probability of a positive result as $\nu((p_0, 1]) = (1 - p_0)\alpha + p_0(1 - \beta)$. Thus, a uniform increase in the rewards which a results-based contract offers for positive results creates an incentive for researchers to perform experiments with a higher rate of false positive results. This is intuitive. However, it also creates an incentive to perform experiments with a *lower* rate of false negative results. In existing models where researchers simply choose how much noise to introduce into the research process, these effects work against each other. The more general model of experiment choice presented here allows them instead to work in concert. Without this generality, it would be impossible to see that an optimal results-based contract must reward negative results — a key lesson of Theorem 7.

However, the assumption that the researcher has access to *all* possible experiments is a double-edged sword with respect to generality. In the real world, the universe of available statistical methods — and thus experiments — is incomplete. I do not consider such constraints here, since I feel that doing so would distract from the analysis and design of researchers' incentives.¹¹

Turning to the form of the practitioner welfare function w , consider the following. When a positive result $p > p_0$ becomes marginally more persuasive, w increases by $\rho\Phi(\rho p)$. This is due to two effects. First, it affects the extensive margin: observing it causes more practitioners to change their action. Second, it affects the intensive margin: those who already would take the risky action are better off in the interim since the result more reliably indicates that $\omega = 1$. But because the marginal practitioner is *by definition* indifferent

¹¹In addition, when these restrictions are “smooth”, they can be described through a suitable choice of the researcher's cost function.

about which action to take, the extensive margin effect is of second-order magnitude! Hence, the important effect is that on the intensive margin — where, after observing p , a population of size $\Phi(\rho p)$ receives ρ marginally more reliably.

Though the institution in this model is altruistic, it is *not* a social planner. It does not seek to maximize the combined welfare of researcher and practitioners, but instead to extract practitioner surplus from the researcher. If its goal were the former, the optimal contract would be simple: “sell the firm to the agent” by awarding the entire surplus to the researcher with a contract $\psi = w$. Such a contract seems unrealistic, and so do the institutional preferences that would generate it. Instead of being run for the benefit of researchers, research institutions generally have an explicit mission to fund the production of research for the benefit of its users.¹²

My strategy for proceeding is as follows. First, I characterize the optimal direct revelation contract in Section 3. Next, I derive sufficient conditions for the implementation of a choice function $\chi : \Theta \rightarrow X$ via a results-based contract in Section 4. Finally, I characterize the results-based contract which implements the choice function from the optimal direct revelation contract, to which it is revenue equivalent, in Section 4.2.

3 Direct Revelation Contracts

The revelation principle ensures that some direct revelation contract (or equivalently, some methods-based contract) is optimal among all contracts. In this section, I solve for that contract, giving the result in Theorem 4. This allows me to reach my main results (Theorem 7) in the next section by finding and characterizing a results-based contract which induces the same outcomes.

Write the indirect utility of a type θ researcher from reporting $\hat{\theta}$ under the direct revelation contract (χ, τ) as

$$U(\theta, \hat{\theta}, \chi, \tau) \equiv u(\chi(\hat{\theta}), \theta, \tau(\hat{\theta})) = \tau(\hat{\theta}) + E_{\chi[\hat{\theta}]}[H(p, \theta) - H(p_0, \theta)].$$

Then the institution’s optimal direct revelation contract solves

$$\max_{\chi, \tau} E_F[E_{\chi[\theta]}[w(p)] - \tau(\theta)] \tag{1}$$

$$\text{s.t. } E_{\chi[\theta]}p = p_0 \qquad \forall \theta \in [0, 1]$$

$$U(\theta, \theta, \chi, \tau) \geq U(\theta, \theta', \chi, \tau) \qquad \forall \theta, \theta' \in [0, 1] \tag{2}$$

$$U(\theta, \theta, \chi, \tau) \geq 0 \qquad \forall \theta \in [0, 1] \tag{3}$$

The sequence of steps I take to solve this problem is standard in the mechanism design literature (see, e.g., Myerson (1981)). However, because the researcher’s choice variable is a distribution instead of a real vector, the execution of these steps is nonstandard. First, I show that incentive compatibility (2)

¹²For instance, the second objective listed in UW-Madison’s mission statement is “Generate new knowledge through a broad array of scholarly, research and creative endeavors, which provide a foundation for dealing with the immediate and long-range needs of society.”

and the participation constraint (3) are together equivalent to the combination of an envelope condition, a monotonicity condition, and voluntary participation by the low type. Next, I substitute the incentive compatibility constraint into the institution's objective. This produces a problem which can be solved by maximizing a value function for each type θ , subject to a monotonicity constraint. Third, I show how to solve the type-by-type optimization problem when the monotonicity constraint is ignored. Finally, I show that under a *regularity* condition analogous to those common in the mechanism design literature, the solution to the unconstrained problem satisfies the monotonicity constraint.

3.1 Simplifying the Problem

Lemma 1. *The direct revelation contract (χ, τ) satisfies incentive compatibility (2) and the participation constraint (3) if and only if*

i. (Envelope Condition)

$$U(\theta, \theta, \chi, \tau) = U(0, 0, \chi, \tau) + \int_0^\theta E_{\chi[r]} [H_\theta(p, r) - H_\theta(p_0, r)] dr \quad \forall \theta \in [0, 1],$$

ii. (Quasi-Monotonicity)

$$\int_{\theta'}^\theta E_{\chi[r]} [H_\theta(p, r) - H_\theta(p_0, r)] dr \geq \int_{\theta'}^\theta E_{\chi[\theta']} [H_\theta(p, r) - H_\theta(p_0, r)] dr \quad \forall \theta, \theta' \in [0, 1],$$

iii. (Low Type Participation)

$$U(0, 0, \chi, \tau) \geq 0.$$

This result is similar to others in the mechanism design literature. However, due to the fact that the researcher chooses a *distribution* rather than a vector, its proof requires the full strength of the techniques of Milgrom and Segal (2002). Note that the resulting monotonicity condition is weaker than the requirement that higher types produce more informative experiments. A corollary to the envelope condition is revenue equivalence:

Corollary 1 (Revenue Equivalence). *If (χ, τ) and $(\chi, \hat{\tau})$ are incentive compatible direct revelation contracts with $V(0, 0, \chi, \tau) = V(0, 0, \chi, \hat{\tau})$, then $\tau = \hat{\tau}$.*

As in many other mechanism design problems, the substitution of the envelope condition into the institution's objective function in a way which allows type-by-type optimization then follows from a clever integration by parts.

Lemma 2. *A direct revelation contract (χ^*, τ^*) solves the optimization problem (1) if and only if χ^* solves*

$$\max_{\chi} E_F \left[E_{\chi[\theta]} \left[w(p) - \frac{1 - F(\theta)}{f(\theta)} (H_\theta(p, \theta) - H_\theta(p_0, \theta)) + H(p, \theta) - H(p_0, \theta) \right] \right] \quad (4)$$

$$s.t. E_{\chi[\theta]} p = p_0, \quad \forall \theta \in [0, 1]$$

$$\int_{\theta'}^\theta E_{\chi[r]} [H_\theta(p, r)] dr \geq \int_{\theta'}^\theta E_{\chi[\theta']} [H_\theta(p, r)] dr, \quad \forall \theta, \theta' \in [0, 1]$$

and τ^* is given by

$$\tau^*(\theta) = \int_0^\theta E_{\chi^*[r]}(H_\theta(p, r) - H_\theta(p_0, r))dr - E_{\chi^*[\theta]}(H(p, \theta) - H(p_0, \theta)).$$

I will consider the problem absent the quasi-monotonicity constraint, and thus find the optimal experiment for each type separately. Then, once I have found the unconstrained problem's solution, I will show that it in fact satisfies the constraint.

Since the choice variable is a distribution rather than a vector, finding this solution is not straightforward. Doing so will require making use of the contributions of Kamenica and Gentzkow (2011). However, I go a step further by using differentiability to characterize the optimal experiment's support as the solution to a system of equations. These results apply to all binary-state Bayesian persuasion problems, and allow me to give comparative statics results using the implicit function theorem. These comparative statics results are new to the literature, and are indispensable to my analysis. In this section, I use them to show that under a regularity condition, the solution to the institution's type-by-type problem is monotone. Later, in section 4.2, I use them to sign the derivatives of the optimal results-based contract, when one exists.

3.2 Optimal Experiments

The institution's type-by-type optimization problem for type θ consists of choosing an experiment $\chi^*[\theta]$ to maximize the expectation of a *value function* $\pi(p, \theta)$ of the posterior beliefs p it induces. That is,

$$\chi^*[\theta] = \arg \max_{\nu} \{E_{\nu} \pi(p, \theta) \text{ s.t. } E_{\nu} p = p_0\}$$

where

$$\pi(p, \theta) = w(p) - \frac{1 - F(\theta)}{f(\theta)} (H_\theta(p, \theta) - H_\theta(p_0, \theta)) + H(p, \theta) - H(p_0, \theta).$$

This problem is identical to the sender's problem in the Bayesian persuasion game described by Kamenica and Gentzkow (2011). Unlike most differentiable settings, solving it is not as simple as using standard first- and second-order conditions for maxima: the choice variable is not a point, but a *distribution*. Instead, I develop analogous necessary and sufficient conditions for solutions to differentiable Bayesian persuasion problems. The scope of these results reaches beyond the setting of my paper: Unless otherwise noted, all of the results in section 3.2 apply to the *general* binary-state Bayesian persuasion problem

$$\max_{\nu} \{E_{\nu} \pi(p, \theta) \text{ s.t. } E_{\nu} p = p_0\}$$

for *arbitrary* π .

Let $\text{conv}(\pi(\cdot, \theta))$ be the convex hull of the graph of $\pi(\cdot, \theta)$. Define

$$\Pi(p, \theta) \equiv \sup\{y \mid (p, y) \in \text{conv}(\pi(\cdot, \theta))\}.$$

For each θ , $\Pi(\cdot, \theta)$ is the smallest concave function which is at least as large as $\pi(\cdot, \theta)$. The following is then straightforward.¹³

¹³Aumann et al. (1995) derive a similar result in the context of repeated zero-sum games of imperfect information.

Lemma 3 (Kamenica and Gentzkow (2011)). *If $\pi(\cdot, \theta)$ is upper semicontinuous, $\arg \max_{\nu} \{E_{\nu} \pi(p, \theta) \text{ s.t. } E_{\nu} p = p_0\}$ is nonempty and $\Pi(p_0, \theta) = \max_{\nu} \{E_{\nu} \pi(p, \theta) \text{ s.t. } E_{\nu} p = p_0\}$.*

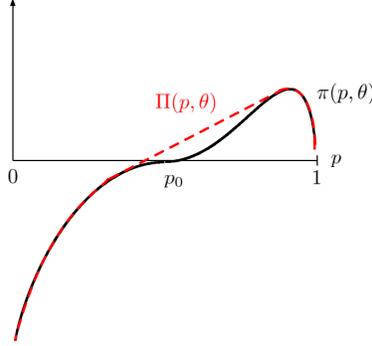


Figure 1: $\pi(\cdot, \theta)$ and its concave closure $\Pi(\cdot, \theta)$. This figure shows a value function $\pi(\cdot, \theta)$ and its concave closure. Note that the value function has two “peaks”, and that the “bridge” which the graph of the concave closure $\Pi(\cdot, \theta)$ creates between them forms part of the tangent line to $\pi(\cdot, \theta)$ at each of its ends. The intuition that this will be true more generally is verified in Theorem 1, and the intuition that this tangent line condition pins down the endpoints of such a “bridge” over a double-peaked value function uniquely is confirmed in Theorem 2.

In my environment, both practitioner welfare and the researcher’s virtual valuation — the sum of which yields the institution’s value function — are continuous. Hence, Lemma 3 ensures that a solution to the institution’s problem exists.

In the following lemma, I refine the characterization of solutions ν^* from Kamenica and Gentzkow (2011) by exploiting the fact that beliefs are real numbers in binary-state Bayesian persuasion problems:

Lemma 4. *Suppose $\pi(\cdot, \theta)$ is upper semicontinuous. $\nu^* \in \arg \max_{\nu} \{E_{\nu} \pi(p, \theta) \text{ s.t. } E_{\nu} p = p_0\}$ if and only if ν^* is Bayes-plausible and for all $p \in \text{supp } \nu^*$,*

- i. $\Pi(p, \theta) = \pi(p, \theta)$, and*
- ii. Either $\Pi(p, \theta) = \Pi(p_0, \theta) + \Pi_p(p_0, \theta)(p - p_0)$ or $p = p_0$.*

For each such p ,

- iii. If $\pi(\cdot, \theta)$ is differentiable at p , then $\pi_p(p, \theta) = \Pi_p(p_0, \theta)$.*
- iv. If $\pi(\cdot, \theta)$ is twice differentiable at p , then $\pi_{pp}(p, \theta) \leq 0$.*

The first part of Lemma 4 tells us that the set of optimal experiments is characterized by Bayes-plausibility and two conditions, (i) and (ii), on their support. The first condition stipulates that the values of the concave closure $\Pi(\cdot, \theta)$ and the value function $\pi(\cdot, \theta)$ coincide on $\text{supp } \nu^*$. The second is equivalent to the requirement that the concave closure $\Pi(\cdot, \theta)$ is an affine function on the convex hull of $\text{supp } \nu^*$. An important implication is that to find $\nu^* \in \arg \max_{\nu} \{E_{\nu} \pi(p, \theta) \text{ s.t. } E_{\nu} p = p_0\}$, we need only find $\text{supp } \nu^*$ which satisfies (i) and

(ii). The second part of the lemma gives us two more necessary conditions on the beliefs which support an optimal experiment. (iii) tells us that at every point in $\text{supp } \nu^*$ at which $\pi(\cdot, \theta)$ is differentiable, $\pi(\cdot, \theta)$ is tangent to its concave closure. Finally, (iv) tells us that in order to form part of the concave closure, the tangent line to $\pi(\cdot, \theta)$ at $p \in \text{supp } \nu^*$ must lie above the value function.

Lemma 4 characterizes the beliefs which support optimal experiments. The following result from Kamenica and Gentzkow (2009) tells us we only need two of these beliefs to pin down a solution $\nu^* \in \arg \max_{\nu} \{E_{\nu} \pi(p, \theta) \text{ s.t. } E_{\nu} p = p_0\}$.

Lemma 5 (Kamenica and Gentzkow (2009)). *There exists $\nu^* \in \arg \max_{\nu} \{E_{\nu} \pi(p, \theta) \text{ s.t. } E_{\nu} p = p_0\}$ such that $\text{supp } \nu^*$ has no more than two elements.*

Together, these results allow me to extend the real-valued maximization concepts of *first-order conditions* and *critical points* to the Bayesian persuasion setting.

Definition (Critical Pair). $\{\underline{p}, \bar{p}\}$ is a *critical pair* of $\pi(\cdot, \theta)$ if $\underline{p} < p_0 < \bar{p}$ and

- if $\pi(\cdot, \theta)$ is differentiable at \underline{p} , $\pi(\underline{p}, \theta) + \pi_p(\underline{p}, \theta)(\bar{p} - \underline{p}) - \pi(\bar{p}, \theta) = 0$;
- if $\pi(\cdot, \theta)$ is differentiable at \bar{p} , $\pi(\bar{p}, \theta) + \pi_p(\bar{p}, \theta)(\bar{p} - \underline{p}) - \pi(\bar{p}, \theta) = 0$;
- if $\pi(\cdot, \theta)$ is twice differentiable at \underline{p} , $\pi_{pp}(\underline{p}, \theta) \leq 0$; and
- if $\pi(\cdot, \theta)$ is twice differentiable at \bar{p} , $\pi_{pp}(\bar{p}, \theta) \leq 0$.

$\{\underline{p}, \bar{p}\}$ is an *interior critical pair* of $\pi(\cdot, \theta)$ if it is a critical pair with $\underline{p} > 0$ and $\bar{p} < 1$.

The conditions

$$\begin{aligned} \pi(\underline{p}, \theta) + \pi_p(\underline{p}, \theta)(\bar{p} - \underline{p}) &= \pi(\bar{p}, \theta), \\ \pi(\bar{p}, \theta) + \pi_p(\bar{p}, \theta)(\bar{p} - \underline{p}) &= \pi(\bar{p}, \theta) \end{aligned}$$

are the Bayesian persuasion analogues to first-order conditions. Instead of looking for *one* point at which the tangent line to the value function is *flat*, here we need to look for *two* points at which the tangent line to the value function is *the same*. They are necessary for a binary-support distribution to be an extremum of the objective in a Bayesian persuasion problem.

Meanwhile, the conditions on the second derivative, when it exists, are the Bayesian persuasion analogues to second-order conditions. Like their cousins in real-valued optimization, they are necessary for an extremum of the objective to be a maximum: the tangent line connecting the critical pair must lie *above*, not below, the graph of the value function.

The above statements are formalized below in Theorem 1, which follows directly from Lemmas 4 and 5.

Theorem 1 (Necessary Conditions on the Set of Maximizers). *Suppose $\pi(\cdot, \theta)$ is upper semicontinuous.*

- i. For any $\nu^* \in \arg \max_{\nu} \{E_{\nu} \pi(p, \theta) \text{ s.t. } E_{\nu} p = p_0\}$, each $\{\underline{p}, \bar{p}\} \subseteq \text{supp } \nu^*$ with $\underline{p} < p_0 < \bar{p}$ is a critical pair of $\pi(\cdot, \theta)$.
- ii. There exists $\nu^* \in \arg \max_{\nu} \{E_{\nu} \pi(p, \theta) \text{ s.t. } E_{\nu} p = p_0\}$ such that $\text{supp } \nu^*$ is either a critical pair of $\pi(\cdot, \theta)$ or the singleton $\{p_0\}$.

We know that there must either be an experiment whose support is a critical pair or a totally uninformative experiment which solves the problem. We also know how to find critical pairs. However, we do not know which critical pair (if any) supports a solution.

In real-valued optimization, we can evaluate the objective function at each critical point to find a global maximum. Similarly, in Bayesian persuasion, we can evaluate the expectation of the value function under each experiment whose support is a critical pair or a singleton at the prior. However, just as a single-peaked (i.e., strictly concave) objective has only a single interior critical point in real-valued optimization problems, a *double-peaked* value function has only a single interior critical pair.

Definition (Double-Peaked). $\pi(\cdot, \theta)$ is *double-peaked* if for $0 < p_1 < p_2 < 1$,

- $\pi(\cdot, \theta)$ is continuously differentiable;
- $\pi(\cdot, \theta)$ is twice continuously differentiable with $\pi_{pp}(p, \theta) < 0$ on $(0, p_1) \cup (p_2, 1)$;
- $\pi(\cdot, \theta)$ is twice continuously differentiable with $\pi_{pp}(p, \theta) > 0$ on (p_1, p_2) ; and
- $\pi(\cdot, \theta)$ satisfies an Inada condition: $\lim_{p \rightarrow 0} \pi_p(p, \theta) = \lim_{p \rightarrow 1} -\pi_p(p, \theta) = \infty$.

Essentially, double-peakedness amounts to the combination of differentiability, an Inada condition, and a single interval of convexity between two intervals of concavity.

Theorem 2 (Sufficiency and Uniqueness for Double-Peaked Problems). *Suppose $\pi(\cdot, \theta)$ is double-peaked.*

- i. $\pi(\cdot, \theta)$ has at most one interior critical pair.
- ii. If $\pi(\cdot, \theta)$ has no interior critical pair, then the totally uninformative experiment ν_0 with $\text{supp } \nu_0 = \{p_0\}$ is the unique solution to $\max_{\nu} \{E_{\nu} \pi(p, \theta) \text{ s.t. } E_{\nu} p = p_0\}$.
- iii. If $\pi(\cdot, \theta)$ has an interior critical pair, then the Bayes-plausible ν^* it supports is the unique solution to $\max_{\nu} \{E_{\nu} \pi(p, \theta) \text{ s.t. } E_{\nu} p = p_0\}$.

Because of the way w and H are formulated in my model, the institution's value function satisfies the differentiability and Inada conditions. Therefore, the bite of double-peakedness in my setting is that to the right of the prior, the convexity of the practitioner welfare function is only greater than the concavity of the virtual valuation on a single interval.

Example 1. A sufficient condition for the institution's value function to have only a single interval of convexity is that $\pi_{ppp}(p, \theta) < 0$ on the interval $(\underline{s}/\rho, \bar{s}/\rho)$ where practitioner welfare is increasing. Suppose once again that the cost of an experiment is proportional to the expected change in the entropy of beliefs it produces:

$$H(p, \theta) = -g(\theta)(p \log(p) + (1-p) \log(1-p)), \quad g(\theta) > 0, g'(\theta) < 0.$$

Then we have

$$\pi_{ppp}(p, \theta) = \left(g(\theta) - \frac{1-F(\theta)}{f(\theta)} g'(\theta) \right) \left(\frac{1}{p^2} - \frac{1}{(1-p)^2} \right) + \rho^3 \phi'(\rho p), \quad \frac{\underline{s}}{\rho} < p < \frac{\bar{s}}{\rho}$$

The first term is negative for all $p > 1/2$. So if, for instance,

- practitioners only change their action when they think the alternative hypothesis $\omega = 1$ is more likely than not to be true (i.e., $\underline{s} > \rho/2$), and
- more practitioners have lower-value outside options than higher-value outside options (i.e., $\phi'(s) \leq 0$),

then the institution's value function is double-peaked for each θ .

When the value function is double-peaked, Theorem 2 tells us that the solution $\chi^*[\theta]$ to the institution's type-by-type optimization problem is the experiment which induces the posteriors given by the interior critical pair of $\pi(\cdot, \theta)$, if it exists, and a totally uninformative experiment, if it does not.

It remains to be shown that this χ^* satisfies the quasi-monotonicity constraint, and so solves the institution's problem from Lemma 2. Lemma 6 and Theorem 3 show that this is the case, so long as the convexity of the value function is nondecreasing in θ . First, Lemma 6 shows that the types θ for whom $\chi^*[\theta]$ is informative are each higher than all types for whom $\chi^*[\theta]$ is uninformative, and that the former set is open. Define

$$\Theta_0 \equiv \{\theta \in \Theta \mid \Pi(p_0, \theta) = 0\}, \quad \Theta_+ \equiv \{\theta \in \Theta \mid \Pi(p_0, \theta) > 0\}, \quad \underline{\theta} \equiv \inf \Theta_+.$$

Lemma 6 (Monotone Participation). *Suppose that for all θ , $\pi(\cdot, \theta)$ is double-peaked and $\pi(p_0, \theta) = 0$.*

- i. *For all $\theta \in \Theta_0$, no interior critical pair of $\pi(\cdot, \theta)$ exists.*
- ii. *For all $\theta \in \Theta_+$, a unique interior critical pair $\{\underline{p}(\theta), \bar{p}(\theta)\}$ of $\pi(\cdot, \theta)$ exists.*
- iii. *If $\pi(p, \theta') - \pi(p, \theta)$ is a convex function of p for each $\theta', \theta \in \Theta$, $\theta' > \theta$, then $\Pi(p_0, \theta)$ is nondecreasing in θ . Hence, either $\Theta_+ = \emptyset$, $\Theta_+ = \Theta$, or $\Theta_+ = (\underline{\theta}, 1]$.*

Second, the interior critical pair of $\pi(\cdot, \theta)$ is getting further from the prior — and thus the experiment it supports is getting more informative — as θ increases. This result is based on the implicit function theorem, and so requires a smoothness condition on π :

Definition (Smoothly Double-Peaked). π is *smoothly double-peaked* if

- for each θ , $\pi(\cdot, \theta)$ is double-peaked with inflection points $p_1(\theta), p_2(\theta)$;
- $\pi_{p\theta}$ exists and is continuous on $(0, 1) \times (0, 1)$; and
- for all $p \neq p_1(\theta), p_2(\theta)$ there is a neighborhood of (p, θ) in which π_{pp} exists and is continuous.

Smooth double-peakedness ensures that at every θ for which $\pi(\cdot, \theta)$ has a critical pair $\{\underline{p}(\theta), \bar{p}(\theta)\}$, there is a neighborhood of $(\theta, \underline{p}(\theta), \bar{p}(\theta))$ on which the system of tangent line equations is continuously differentiable. This is exactly what we need for the implicit function theorem to be applicable. Requiring continuity of $\pi_{p\theta}$ allows us to sign $\underline{p}'(\theta)$ and $\bar{p}'(\theta)$ using the fact that concavity is increasing in θ . In my setting, the value function is smoothly double-peaked if and only if it is double-peaked for each θ .

Theorem 3 (Convexity Comparative Statics). *Suppose π is smoothly double-peaked and that $\pi(p_0, \theta) = 0$ for all θ .*

- If the convexity of $\pi(\cdot, \theta)$ is nondecreasing in θ (i.e., if $\pi_{p\theta}$ is nondecreasing in p), then \underline{p} is nonincreasing, and \bar{p} nondecreasing, in θ on Θ_+ .*
- If the convexity of $\pi(\cdot, \theta)$ is increasing in θ (i.e., if $\pi_{p\theta}$ is increasing in p), then \underline{p} is decreasing, and \bar{p} increasing, in θ on Θ_+ .*

Note that in my setting, the practitioner welfare function depends only on the experiment's result, not on the researcher's type. Therefore, the convexity of the institution's value function is increasing in θ if and only if the convexity of the researcher's virtual valuation is. A condition analogous to the regularity conditions common in the mechanism design literature ensures that this is the case.

Definition (Regularity). The type distribution is *regular* if $\frac{\partial}{\partial \theta} \left(H_{pp}(\theta, p) - \frac{1-F(\theta)}{f(\theta)} H_{\theta pp}(\theta, p) \right) \geq 0$ for all $\theta, p \in (0, 1)$ and *strictly regular* if $\frac{\partial}{\partial \theta} \left(H_{pp}(\theta, p) - \frac{1-F(\theta)}{f(\theta)} H_{\theta pp}(\theta, p) \right) > 0$ for all $\theta, p \in (0, 1)$.

Example 2. Suppose yet again that the cost of an experiment is proportional to the expected change it produces in the entropy of beliefs. In this context, regularity means that

$$- \left(g'(\theta) \left(1 - \frac{d}{d\theta} \left[\frac{1-F(\theta)}{f(\theta)} \right] \right) - \frac{1-F(\theta)}{f(\theta)} g''(\theta) \right) \left(\frac{1}{p} + \frac{1}{1-p} \right) \geq 0$$

If $g''(\theta) \geq 0$, this amounts to the requirement of a monotone hazard rate.¹⁴

Theorem 4 (Optimal Direct Revelation Contract). *If the type distribution is regular and the institution's value function*

$$\pi(p, \theta) = w(p) - \frac{1-F(\theta)}{f(\theta)} (H_\theta(p, \theta) - H_\theta(p_0, \theta)) + H(p, \theta) - H(p_0, \theta)$$

is double-peaked for all $\theta \in \Theta$, then the institution's unique optimal direct revelation contract is given by the experiment choice function χ^ such that*

$$\text{supp } \chi^*[\theta] = \begin{cases} \{\underline{p}(\theta), \bar{p}(\theta)\}, & \theta \in \Theta_+, \\ \{p_0\}, & \theta \in \Theta_0, \end{cases}$$

¹⁴In fact, this is true whenever H is multiplicatively separable.

where for each $\theta \in \Theta_+$, $\{\underline{p}(\theta), \bar{p}(\theta)\}$ is the unique critical pair of $\pi(\cdot, \theta)$, and the transfer function

$$\tau^*(\theta) = \int_0^\theta E_{\chi^*[\theta]}(H_\theta(p, r) - H_\theta(p_0, r))dr - E_{\chi^*[\theta]}(H(p, \theta) - H(p_0, \theta)), \quad \theta \in \Theta.$$

The main purpose of solving for this contract is so that I can ask whether a results-based contract can be optimal among all contracts. However, it also allows me to consider the way that the optimal contract distorts outcomes away from efficiency. In this regard, the institution's optimal contract from Theorem 4 exhibits standard features of optimal contracts in environments with adverse selection. Namely, the outcome is distorted away from efficiency for all but the highest type. In my environment, this occurs because the virtual valuation is more concave than the measure of uncertainty H that determines researcher costs. Hence, researchers undertake less informative experiments (in the Blackwell sense) than would be efficient.

It is also interesting to note that in the limit as $\theta \rightarrow \underline{\theta} \equiv \inf \Theta_+$, it is impossible for *both* elements of the critical pair to near the prior p_0 , because there must be a gap between the inflection points $p_1(\theta), p_2(\theta)$ which separate them, which cannot get arbitrarily small.

Lemma 7. *Suppose that π is smoothly double-peaked. For $\theta \in \Theta_+$, let $\{\underline{p}(\theta), \bar{p}(\theta)\}$ be the unique critical pair of $\pi(\cdot, \theta)$. $\lim_{\theta \rightarrow \underline{\theta}} \bar{p}(\theta) - \underline{p}(\theta) > 0$.*

Lemma 7 also ensures that the results-based implementation techniques in Section 4 are applicable to χ^* .

4 Results-based Contracts

The problem of contracting on results is different from that of contracting on methods (or using a direct revelation contract) in three substantive ways.

First, payment occurs after the result of a researcher's experiment has been realized. However, the researcher is risk-neutral, so a results-based contract ψ is outcome-equivalent to the methods-based contract with $T(\nu) = E_\nu[\psi(p)]$ for all ν .

This ensures that a results-based contract can never give the institution a higher payoff than the optimal direct revelation contract (χ^*, τ^*) found in Theorem 4. Moreover, due to revenue equivalence (Corollary 1), it achieves the same institutional payoffs if and only if it causes each type θ to choose the experiment $\chi^*[\theta]$, and yields zero payoff for type 0 researchers.

It also implies the second difference: in a results-based contract, there must be some function ψ such that the expected transfer for each ν is $E_\nu[\psi(p)]$. This costs the institution a degree of freedom in contracting. Finally, the researcher need not reveal a result if doing so would be costly. This means that the participation constraint is ex post: not only must $E_\nu[\psi(p)] \geq 0$ for each ν , but $\psi(p) \geq 0$ for each p .

This section investigates whether these factors constrain the institution when contracting on results instead of methods. Under quite general conditions, the answer is no. Further, it provides qualitative results on the shape of the optimal results-based contract.

4.1 Implementation

When facing the results-based contract ψ , the researcher's optimization problem is

$$\max_{\nu} \{E_{\nu} v(p, \theta, \psi) \text{ s.t. } E_{\nu} p = p_0\},$$

where

$$v(p, \theta, \psi) = \psi(p) + H(p, \theta) - H(p_0, \theta).$$

Similarly to the pointwise formulation of the direct revelation contract design problem, this problem is one of Bayesian persuasion. However, instead of requiring us to find a solution given a value function, the contract design problem is the reverse: we must find a reward function $\psi : [0, 1] \rightarrow \mathbb{R}_+$ given solutions $\chi : \Theta \rightarrow X$.

In doing so, it will be helpful to restrict attention to a class of χ which includes the χ^* found in Theorem 4.

Definition (\underline{y}, \bar{y} -Binary, \underline{y}, \bar{y} -Smoothness, Strict Monotonicity of χ). $\chi : \Theta \rightarrow X$ is \underline{y}, \bar{y} -binary if for $\tilde{\theta} \in \Theta$, $\underline{y} : [\tilde{\theta}, 1] \rightarrow [0, p_0]$, and $\bar{y} : [\tilde{\theta}, 1] \rightarrow [p_0, 1]$ with $\underline{y}(\theta) \neq \bar{y}(\theta) \forall \theta$, either

$$\text{supp } \chi[\theta] = \begin{cases} \{\underline{y}(\theta), \bar{y}(\theta)\}, & \theta \geq \tilde{\theta}, \\ \{p_0\}, & \theta < \tilde{\theta}; \end{cases} \quad \text{or} \quad \text{supp } \chi[\theta] = \begin{cases} \{\underline{y}(\theta), \bar{y}(\theta)\}, & \theta > \tilde{\theta}, \\ \{p_0\}, & \theta \leq \tilde{\theta}. \end{cases}$$

It is *strictly monotone* if $\chi[\theta]$ is a mean-preserving spread of $\chi[\theta']$ for all $\theta > \theta'$. It is \underline{y}, \bar{y} -smooth if it is \underline{y}, \bar{y} -binary and \underline{y}, \bar{y} are continuously differentiable.

Theorem 1 tells us that if a results-based contract ψ implements a \underline{y}, \bar{y} -binary choice function χ , then for each θ for which $\text{supp } \chi[\theta] = \{\underline{y}(\theta), \bar{y}(\theta)\}$, $\{\underline{y}(\theta), \bar{y}(\theta)\}$ must be a critical pair of $v(\cdot, \theta, \psi)$. Therefore, ψ must solve

$$\begin{aligned} v(\underline{y}(\theta), \theta, \psi) + v_p(\underline{y}(\theta), \theta, \psi)(\bar{y}(\theta) - \underline{y}(\theta)) - v(\bar{y}(\theta), \theta, \psi) &= 0, \\ v(\underline{y}(\theta), \theta, \psi) + v_p(\bar{y}(\theta), \theta, \psi)(\bar{y}(\theta) - \underline{y}(\theta)) - v(\bar{y}(\theta), \theta, \psi) &= 0 \end{aligned}$$

for each such θ .

When ψ makes $v(\cdot, \theta, \psi)$ double-peaked, this alone is sufficient for $\chi[\theta]$ to solve the type- θ researcher's problem when facing the results-based contract ψ . Unfortunately, this will not generally be the case. Instead, we need a weaker sufficient condition for a critical pair to support a global maximum of the objective in a Bayesian persuasion problem.

Theorem 5. If $\{\underline{p}, \bar{p}\}$ is a critical pair of $v(\cdot, \theta, \psi)$ and

$$v(\underline{p}, \theta, \psi) + \frac{v(\bar{p}, \theta, \psi) - v(\underline{p}, \theta, \psi)}{\bar{p} - \underline{p}}(p - \underline{p}) \geq v(p, \theta, \psi)$$

for all $p \in [0, 1]$, then the Bayes-plausible ν^* it supports solves $\max_{\nu} \{E_{\nu} v(p, \theta, \psi) \text{ s.t. } E_{\nu} p = p_0\}$.

Geometrically, if the graph of $v(\cdot, \theta, \psi)$ lies below a line, then so do its secant lines.

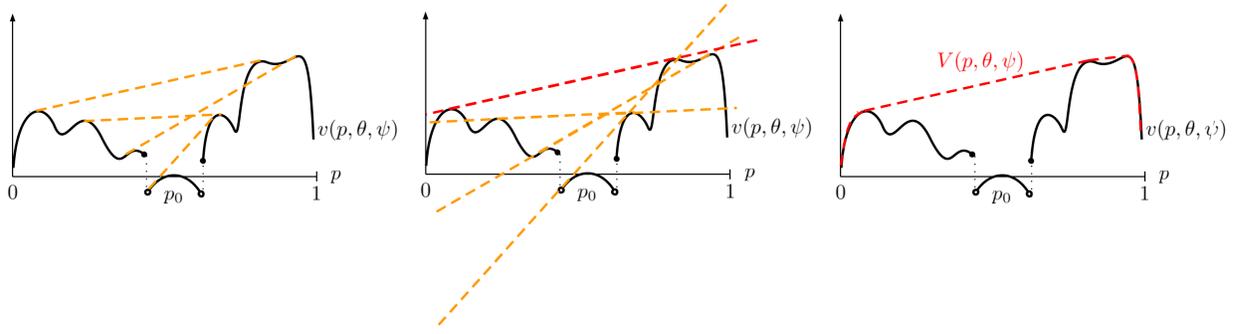


Figure 2: Theorem 5's sufficient condition. **Panel 1:** A value function which is not double-peaked may have many interior critical pairs. **Panel 2:** Only one of the secants connecting these critical pairs — the highest one — lies above the entire value function. **Panel 3:** The concave closure of $v(\cdot, \theta)$.

Theorems 1 and 5 can then be turned into a formal statement giving sufficient conditions for ψ to be a results-based contract which implements χ .

Lemma 8. *Suppose that $\chi : [0, 1] \rightarrow X$ is strictly monotone and \underline{y}, \bar{y} -binary. $\psi : [0, 1] \rightarrow \mathbb{R}$ is a results-based contract which implements χ if ψ is differentiable on $[0, \underline{y}(\tilde{\theta})]$ and $(\bar{y}(\tilde{\theta}), 1]$, upper semicontinuous, and nonnegative, and for all $\theta \geq \tilde{\theta}$,*

- the secant line between $\underline{y}(\theta)$ and $\bar{y}(\theta)$ extends above the graph of $v(\cdot, \theta, \psi)$:

$$\psi(\underline{y}(\theta)) + H(\underline{y}(\theta), \theta) + (H_p(\underline{y}(\theta), \theta) + \psi'(\underline{y}(\theta)))(p - \underline{y}(\theta)) \geq \psi(p) + H(p, \theta) \forall p \in [0, 1]; \quad (5)$$

- $\{\underline{y}(\theta), \bar{y}(\theta)\}$ solves the system of tangent line conditions for $v(\cdot, \theta, \psi)$:

$$\psi(\underline{y}(\theta)) + H(\underline{y}(\theta), \theta) + (H_p(\bar{y}(\theta), \theta) + \psi'(\bar{y}(\theta)))(\bar{y}(\theta) - \underline{y}(\theta)) = \psi(\bar{y}(\theta)) + H(\bar{y}(\theta), \theta), \quad (6)$$

$$\psi(\bar{y}(\theta)) + H(\bar{y}(\theta), \theta) + (H_p(\underline{y}(\theta), \theta) + \psi'(\underline{y}(\theta)))(\bar{y}(\theta) - \underline{y}(\theta)) = \psi(\underline{y}(\theta)) + H(\underline{y}(\theta), \theta); \quad (7)$$

- and the low type $\tilde{\theta}$'s participation constraint binds:

$$\frac{p_0 - \underline{y}(\tilde{\theta})}{\bar{y}(\tilde{\theta}) - \underline{y}(\tilde{\theta})} \left(\psi(\bar{y}(\tilde{\theta})) + H(\bar{y}(\tilde{\theta}), \tilde{\theta}) \right) + \frac{\bar{y}(\tilde{\theta}) - p_0}{\bar{y}(\tilde{\theta}) - \underline{y}(\tilde{\theta})} \left(\psi(\underline{y}(\tilde{\theta})) + H(\underline{y}(\tilde{\theta}), \tilde{\theta}) \right) = H(p_0, \tilde{\theta}); \quad (8)$$

with $\psi'(\underline{y}(\tilde{\theta}))$ and $\psi'(\bar{y}(\tilde{\theta}))$ denoting left and right derivatives, respectively.

Additionally, if (χ, τ) is an incentive compatible direct revelation contract with $V(0, 0, \chi, \tau) = 0$, then $\tau(\theta) = E_{\chi[\theta]}[\psi(p)]$ for all θ .

Equations (6) and (7) specify a system of differential equations which we would like to solve for ψ . (8) gives an initial condition of that system. Applying (5) at p_0 and noting that (8) implies it must bind there gives us another. However, this is a nonstandard problem: the function ψ we are trying to find does not

take the variable θ as an argument directly, but through the functions \underline{y} and \bar{y} . Nevertheless, a solution can be found: since \underline{y} and \bar{y} are monotone, for each p , $\psi(p)$ must satisfy the system (6),(7) for at most one θ .

For any strictly monotone \underline{y}, \bar{y} -smooth $\chi : \Theta \rightarrow X$, define $\psi_\chi : [0, 1] \rightarrow \mathbb{R}$ as follows:

$$\begin{aligned} \psi_\chi(\bar{y}(\theta)) &\equiv H(p_0, \tilde{\theta}) - H(\bar{y}(\tilde{\theta}), \tilde{\theta}) + (\bar{y}(\theta) - p_0)H_p(p_0, \tilde{\theta}) \\ &\quad + \int_{\tilde{\theta}}^{\theta} \frac{H_\theta(\bar{y}(t), t) - H_\theta(\underline{y}(t), t)}{\bar{y}(t) - \underline{y}(t)} (\bar{y}(\theta) - \bar{y}(t)) - H_p(\bar{y}(t), t) \bar{y}'(t) dt, & \forall \theta \geq \tilde{\theta}, \\ \psi_\chi(\underline{y}(\theta)) &\equiv H(p_0, \tilde{\theta}) - H(\underline{y}(\tilde{\theta}), \tilde{\theta}) + (\underline{y}(\theta) - p_0)H_p(p_0, \tilde{\theta}) \\ &\quad + \int_{\tilde{\theta}}^{\theta} \frac{H_\theta(\bar{y}(t), t) - H_\theta(\underline{y}(t), t)}{\bar{y}(t) - \underline{y}(t)} (\underline{y}(\theta) - \underline{y}(t)) - H_p(\underline{y}(t), t) \underline{y}'(t) dt, & \forall \theta \geq \tilde{\theta}, \\ \psi_\chi(p) &\equiv 0, & p \in (\underline{y}(\tilde{\theta}), \bar{y}(\tilde{\theta})), \\ \psi_\chi(p) &\equiv \psi(\bar{y}(1)) + \psi'(\bar{y}(1))(p - \bar{y}(1)), & p > \bar{y}(1), \\ \psi_\chi(p) &\equiv \psi(\underline{y}(1)) + \psi'(\underline{y}(1))(p - \underline{y}(1)), & p < \underline{y}(1). \end{aligned}$$

ψ_χ solves the system of differential equations (6) and (7) for all θ . However, in order to ensure that (5) is satisfied, the cost function must be multiplicatively separable:

Definition (Multiplicative Separability). H is *multiplicatively separable* if there exist $h : [0, 1] \rightarrow \mathbb{R}, g : \Theta \rightarrow \mathbb{R}_+$ such that $H(p, \theta) = h(p)g(\theta)$ for all p, θ .

Note that $H_{pp}(p, \theta) < 0$ and $H_{pp\theta}(p, \theta) > 0$ imply $h''(p) < 0$ and $g'(\theta) < 0$. One especially notable multiplicatively separable measure of uncertainty is the entropy function I use in my examples.

Without multiplicative separability, the concavity of H is changing at different rates for different results p . When this happens, *local* incentive compatibility conditions (the tangent line equations (6) and (7)) are insufficient for the contract to be *globally* incentive compatible (i.e., satisfy the sufficient condition (5)).

Theorem 6 (Implementation Using a Results-Based Contract). *Suppose $\chi : \Theta \rightarrow X$ is strictly monotone and \underline{y}, \bar{y} -smooth, and that H is multiplicatively separable. If ψ_χ is nonnegative, then it is a results-based contract which implements the experiment choice function χ . If, in addition, (χ, τ) is an incentive compatible direct revelation contract with $V(0, 0, \chi, \tau) = 0$, then $\tau(\theta) = E_{\chi[\theta]}[\psi_\chi(p)]$ for all θ .*

To see why χ can be implemented by a results-based contract, consider the following. The results-based implementation problem transforms the envelope incentive compatibility condition for direct revelation contracts to the local incentive compatibility condition given by the system of tangent line equations (6) and (7). Incentive compatibility requires direct revelation contracts to satisfy a weak monotonicity condition; similarly, χ must be monotone for the system (6),(7) to map out the payment function ψ . Finally, the participation constraint (8) must be satisfied.

The main difference between results-based contracting and methods-based contracting is that the institution can no longer refuse payment completely for experiments off the equilibrium path (that is, in $X \setminus \chi[\Theta]$). This adds the new global incentive compatibility condition (5). This global incentive compatibility condition

only impedes implementation if the marginal cost of informativeness (i.e., the convexity of H) is increasing faster in some areas than others. If it does, payments for the results of $\chi[\theta]$ may create a peak in the graph of $v(\cdot, \theta', \psi_\chi)$ for some type $\theta' > \theta$. If this peak rises above the secant line connecting $(\underline{y}(\theta'), v(\underline{y}(\theta'), \theta', \psi_\chi))$ and $(\bar{y}(\theta'), v(\bar{y}(\theta'), \theta', \psi_\chi))$, then $\chi[\theta']$ is no longer optimal for θ' . This is ruled out by multiplicative separability.

Applying Theorem 6 to the experiment choice function χ^* from the optimal direct revelation contract yields the following corollary.

Corollary 2. *If (χ^*, τ^*) is the practitioner-optimal direct revelation contract, then ψ_{χ^*} is a results-based contract which implements χ^* whenever*

1. *the type distribution is strictly regular;*
2. *the institution's value function $\pi(\cdot, \theta)$ is double-peaked for each θ ; and*
3. *$\psi_{\chi^*}(p) \geq 0$ for all $p \in [0, 1]$.*

In this case, the expected costs of both contracts are the same: $\tau^(\theta) = E_{\chi^*[\theta]}[\psi_{\chi^*}(p)]$, and hence ψ_{χ^*} is the institution's optimal results-based contract ψ^* .*

4.2 The Optimal Results-Based Contract

Under a stronger version of regularity, it is possible to show that ψ_{χ^*} is nonnegative. Thus, the ex post participation constraint inherent in results-based contracting does not bind. This condition also allows me to describe the shape of ψ_{χ^*} by signing its derivatives.

Definition (Superregularity). The type distribution is *superregular* if

$$(g'(\theta))^2 \left(\frac{1 - F(\theta)}{f(\theta)} \right) - g(\theta) \frac{d}{d\theta} \left[\frac{1 - F(\theta)}{f(\theta)} g'(\theta) \right] \leq 0.$$

Lemma 9. *If the type distribution is superregular, then it is strictly regular.*

Superregularity is satisfied, for instance, if the hazard rate of F is monotone and $g''(\theta)g(\theta) \geq (g'(\theta))^2$. One g which satisfies the latter condition is $g(\theta) = \frac{1}{a+b\theta}$ for constants $a, b > 0$.

When the type distribution is superregular and the institution's value function is double-peaked, an optimal results-based contract rewards sufficiently informative results, both negative and positive; rewards more informative results more; and does so at a rate which is increasing in informativeness.

Theorem 7 (Characterization of the Optimal Results-Based Contract). *Suppose that H is multiplicatively separable, the type distribution is superregular, the institution's value function $\pi(\cdot, \theta)$ is double-peaked for each θ , and $\Theta_+ \neq \emptyset$. Then*

- i. *(χ^*, τ^*) is the optimal direct revelation contract.*

- ii. ψ_{χ^*} is a results-based contract which implements χ^* , and the expected costs of ψ_{χ^*} and (χ^*, τ^*) are the same: $E_{\chi^*[\theta]}[\psi_{\chi^*}(p)] = \tau^*(\theta)$.
- iii. $\psi_{\chi^*}(p) \geq 0$ for all p .
- iv. $\psi'_{\chi^*}(p) < 0$ for $p < \underline{p}(\underline{\theta})$ and $\psi'_{\chi^*}(p) > 0$ for $p > \bar{p}(\underline{\theta})$.
- v. $\psi''_{\chi^*}(p) \geq 0$ for $p < \underline{p}(\underline{\theta})$ and $p > \bar{p}(\underline{\theta})$.

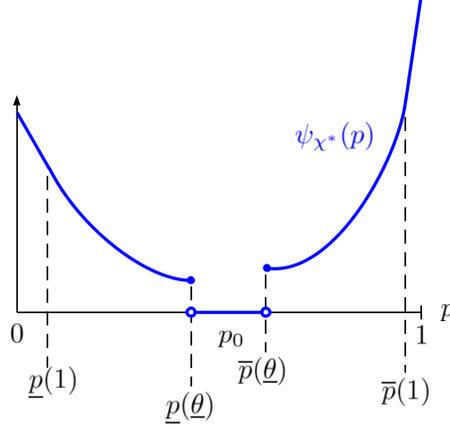


Figure 3: The shape of the optimal results-based contract ψ_{χ^*} . ψ_{χ^*} is convex and decreasing on $[0, \underline{p}(\underline{\theta})]$, zero on $(\underline{p}(\underline{\theta}), \bar{p}(\underline{\theta}))$, and convex and increasing on $[\bar{p}(\underline{\theta}), 1]$.

Theorem 7 makes four important points. First, the optimal direct revelation contract can be implemented with a results-based contract. Second, this results-based contract will reward negative results as well as positive ones, even though the former are not useful to the practitioners. Third, the further that a result moves practitioners' posterior belief away from the prior, the more the contract will reward the researcher. Finally, the contract will reward stronger evidence at an increasing rate.

Example 3. A numerical example is useful in understanding what ψ_{χ^*} looks like in practice. Suppose that

- Researcher costs are proportional to expected change in entropy: $H(p, \theta) = g(\theta)h(p)$ with $h(p) = -(p \log p + (1-p) \log(1-p))$, $g(\theta) = 3/(1+\theta)$.
- Researcher types follow a generalized Pareto distribution on $[0, 1]$ with standard deviation $\sigma = 1/2$: $F(\theta) = 1 - (1-\theta)^2$.
- $\rho = 10$, $\underline{s} = 5$, $\bar{s} = 10$. Thus, practitioners' decision thresholds are distributed between $\frac{1}{2}$ and 1.
- Practitioners' outside options follow a truncated Pareto distribution with parameter $\alpha = 30$: $\Phi(s) = \frac{1 - (\underline{s}/s)^{30}}{1 - (\underline{s}/\bar{s})^{30}}$.

Figure 4 shows the parts of ψ_{χ^*} which contain the support of the experiments $\chi^*[\Theta]$ that researchers actually perform, for three different values of p_0 :

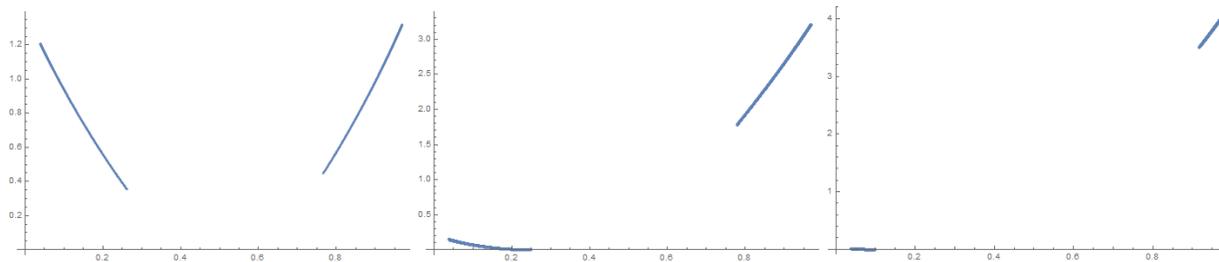


Figure 4: The optimal results-based contract ψ_{χ^*} . Panel 1: $p_0 = 1/2$. Panel 2: $p_0 = 1/4$. Panel 3: $p_0 = 1/10$.

Note that the scales on these graphs are different. This example allows us to see what ψ_{χ^*} looks like and how it changes when the prior shifts. As p_0 moves leftward, experiments must change beliefs more in order to provide value to the practitioners. This causes the optimal contract to implement experiments with a *lower probability* of producing positive results that have a *greater effect* on beliefs. It does so by offering *larger* rewards for positive results and *smaller* rewards for negative results. Figure 5 also causes fewer types to do research, and yields far smaller information rents for those that do:

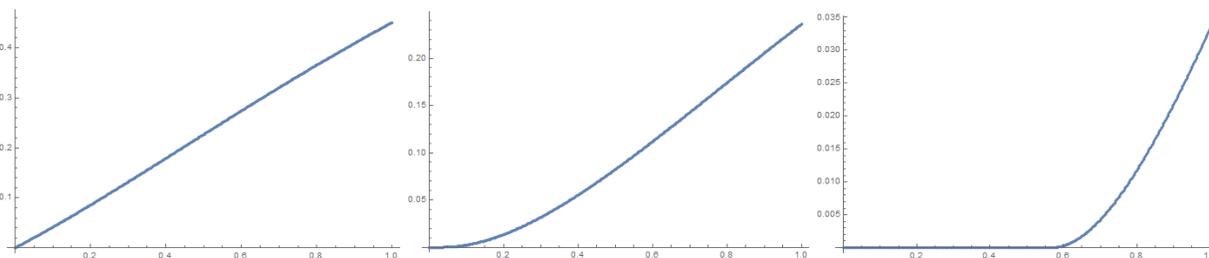


Figure 5: Information rent in the optimal contract as a function of θ . Panel 1: $p_0 = 1/2$. Panel 2: $p_0 = 1/4$. Panel 3: $p_0 = 1/10$. Note that the scales on these graphs differ.

5 Discussion

What have we learned from these results? First, the main reform that Theorem 7 suggests may be worthwhile for institutions is to reward researchers for informative negative results. This is true even when they are not useful to practitioners because they simply offer evidence confirming conventional wisdom. The reason is intuitive. While rewarding negative results encourages experiments which produce false negative results (type II errors) at a higher rate, it also encourages experiments which produce false positive results (type I errors) at a lower rate.¹⁵ Theorem 7 suggests that the latter effect is strong enough to justify rewarding strong evidence in favor of the status quo.

¹⁵For more on type I and II errors, see section 2.1.

Second, it is not automatically true that the outcome of a methods-based contract (or a direct revelation contract) can be implemented by a contract in which reward instead depends on experimental results. However, the institution might still contract on results rather than methods because publication quality only measures the former¹⁶ and is contractible. If this is the case, “results-blind” peer review (i.e., making publication decisions based on methods rather than results) may, in some cases, allow the design of better researcher incentives than results-based peer review. However, if higher researcher quality decreases the cost of each experiment by the same proportion, and the type distribution is superregular, there is no trade-off here. Instead, Theorem 7 shows that there is a results-based contract which is equivalent to the optimal direct revelation contract. The optimal results-based contract then rewards researchers more when they produce stronger evidence — a nontrivial conclusion! Moreover, this is true even when that evidence is in the form of a negative result.

Finally, many researchers question the wisdom of the way we currently reward “top 5” publications. One reason is that the marginal reward to publishing in one of these top journals instead of a less prestigious outlet is disproportionate to the increase in the strength of evidence necessary to do so. However, Theorem 7 shows that when we can characterize it, the optimal results-based contract is convex. This suggests that in the presence of adverse selection, marginal rewards which increase in informativeness may be a feature of the current system, not a bug.

References

- AUMANN, R. J., M. MASCHLER, AND R. E. STEARNS (1995): *Repeated Games with Incomplete Information*, MIT Press.
- BAEZ, J. C., T. FRITZ, AND T. LEINSTER (2011): “A Characterization of Entropy in Terms of Information Loss,” *Entropy*, 13, 1945–1957.
- BERGEMANN, D., A. BONATTI, AND A. SMOLIN (2016): “The Design and Price of Information,” .
- CAPLIN, A. AND M. DEAN (2013): “Behavioral Implications of Rational Inattention with Shannon Entropy,” Working Paper 19318, National Bureau of Economic Research.
- CARROLL, G. (2013): “Robust Incentives for Information Acquisition,” in *Proceedings of the Fourteenth ACM Conference on Electronic Commerce*, ACM, 161–162.
- DI TILLIO, A., M. OTTAVIANI, AND P. N. SORENSEN (2016): “Persuasion Bias in Science: Can Economics Help?” .
- GENTZKOW, M. AND E. KAMENICA (2014): “Costly Persuasion,” *American Economic Review: Papers and Proceedings*, 104, 457–462.

¹⁶Though it does so imperfectly.

- IOANNIDIS, J. P. (2005): “Why Most Published Research Findings are False,” *PLoS Medicine*, 2, e124.
- KAMENICA, E. AND M. GENTZKOW (2009): “Bayesian Persuasion,” Working Paper 15540, National Bureau of Economic Research.
- (2011): “Bayesian Persuasion,” *American Economic Review*, 101.
- (2016): “Competition in Persuasion,” *Review of Economic Studies*, forthcoming.
- KOLOTILIN, A. (2015): “Experimental Design to Persuade,” *Games and Economic Behavior*, 90, 215–226.
- KOLOTILIN, A., M. LI, T. MYLOVANOV, AND A. ZAPECHELNYUK (2015): “Persuasion of a Privately Informed Receiver,” Unpublished Manuscript.
- LIBGOBER, J. (2015): “False Positives in Scientific Research,” *Available at SSRN 2617130*.
- MANIADIS, Z., F. TUFANO, AND J. A. LIST (2014): “One Swallow Doesn’t Make a Summer: New Evidence on Anchoring Effects,” *American Economic Review*, 104, 277–290.
- MILGROM, P. AND I. SEGAL (2002): “Envelope Theorems for Arbitrary Choice Sets,” *Econometrica*, 70, 583–601.
- MYERSON, R. B. (1981): “Optimal Auction Design,” *Mathematics of Operations Research*, 6, 58–73.
- OPEN SCIENCE COLLABORATION (2015): “Estimating the Reproducibility of Psychological Science,” *Science*, 349, aac4716.
- OSBAND, K. (1989): “Optimal Forecasting Incentives,” *Journal of Political Economy*, 97, 1091–1112.
- SHANNON, C. E. (1948): “A Mathematical Theory of Communication,” *Bell System Technical Journal*, 27, 379–423.
- SIMMONS, J. P., L. D. NELSON, AND U. SIMONSOHN (2011): “False-Positive Psychology: Undisclosed Flexibility in Data Collection and Analysis Allows Presenting Anything as Significant,” *Psychological Science*, 22, 1359–1366.
- SMALDINO, P. E. AND R. MCELREATH (2016): “The Natural Selection of Bad Science,” *arXiv preprint arXiv:1605.09511*.
- WASSERSTEIN, R. L. AND N. A. LAZAR (2016): “The ASA’s Statement on p -values: Context, Process, and Purpose,” *The American Statistician*.
- ZERMEÑO, L. (2011): “A Principal-Expert Model and the Value of Menus,” *unpublished paper, Massachusetts Institute of Technology*.

Appendix: Proofs

Proof of Lemma 1 (Envelope, Quasi-Monotonicity, and Low Type Participation Conditions for Incentive Compatibility) (Only if) By the fundamental theorem of calculus, we can write

$$V(\theta, \theta', \chi, \tau) = V(\theta', \theta', \chi, \tau) + E_{\chi[\theta']} \int_{\theta'}^{\theta} H_{\theta}(p, r) - H_{\theta}(p_0, r) dr.$$

Then incentive compatibility (2) can be reformulated as

$$V(\theta, \theta, \chi, \tau) - V(\theta', \theta', \chi, \tau) \geq E_{\chi[\theta']} \int_{\theta'}^{\theta} H_{\theta}(p, r) - H_{\theta}(p_0, r) dr. \quad (9)$$

Suppose (χ, τ) satisfies (2) and (3). Since $H_{\theta pp} > 0$, the integrand is positive and $V(\theta, \theta, \chi, \tau)$ is nondecreasing in θ ; hence, (3) reduces to (iii). Now for any $\theta, \theta' \in [0, 1]$ we have (letting $\theta > \theta'$ without loss)

$$\begin{aligned} |V(\theta, \theta, \chi, \tau) - V(\theta', \theta', \chi, \tau)| &= V(\theta, \theta, \chi, \tau) - V(\theta', \theta', \chi, \tau) \leq E_{\chi[\theta]} \int_{\theta'}^{\theta} H_{\theta}(p, r) - H_{\theta}(p_0, r) dr \\ &\leq \sup_{p \in [0, 1]} \int_{\theta'}^{\theta} H_{\theta}(p, r) - H_{\theta}(p_0, r) dr \\ &\leq \sup_{r, p \in [0, 1]} (H_{\theta}(p, r) - H_{\theta}(p_0, r))(\theta - \theta'), \end{aligned}$$

and so $V(\theta, \theta, \chi, \tau)$ is Lipschitz with constant $\sup_{r, p \in [0, 1]} (H_{\theta}(p, r) - H_{\theta}(p_0, r))$. Therefore, it is differentiable almost everywhere: for almost every θ ,

$$\lim_{\delta \rightarrow 0} \frac{V(\theta + \delta, \theta + \delta, \chi, \tau) - V(\theta, \theta, \chi, \tau)}{\delta} = \lim_{\delta \rightarrow 0} \frac{V(\theta, \theta, \chi, \tau) - V(\theta - \delta, \theta - \delta, \chi, \tau)}{\delta} = \frac{d}{d\theta} V(\theta, \theta, \chi, \tau).$$

Now for all $\theta \in (0, 1)$ and all $0 < \delta < \min\{\theta, 1 - \theta\}$ we have from (9) that

$$\begin{aligned} E_{\chi[\theta]} \frac{\int_{\theta}^{\theta + \delta} H_{\theta}(p, r) - H_{\theta}(p_0, r) dr}{\delta} &\leq \frac{V(\theta + \delta, \theta + \delta, \chi, \tau) - V(\theta, \theta, \chi, \tau)}{\delta}, \\ E_{\chi[\theta]} \frac{\int_{\theta - \delta}^{\theta} H_{\theta}(p, r) - H_{\theta}(p_0, r) dr}{\delta} &\geq \frac{V(\theta, \theta, \chi, \tau) - V(\theta - \delta, \theta - \delta, \chi, \tau)}{\delta}. \end{aligned}$$

Hence,

$$\begin{aligned} \lim_{\delta \rightarrow 0} E_{\chi[\theta]} \frac{\int_{\theta}^{\theta + \delta} H_{\theta}(p, r) - H_{\theta}(p_0, r) dr}{\delta} &\leq \liminf_{\delta \rightarrow 0} \frac{V(\theta + \delta, \theta + \delta, \chi, \tau) - V(\theta, \theta, \chi, \tau)}{\delta}, \\ \lim_{\delta \rightarrow 0} E_{\chi[\theta]} \frac{\int_{\theta - \delta}^{\theta} H_{\theta}(p, r) - H_{\theta}(p_0, r) dr}{\delta} &\geq \limsup_{\delta \rightarrow 0} \frac{V(\theta, \theta, \chi, \tau) - V(\theta - \delta, \theta - \delta, \chi, \tau)}{\delta}. \end{aligned}$$

So for almost every θ ,

$$\begin{aligned} E_{\chi[\theta]} H_{\theta}(p, \theta) - H_{\theta}(p_0, \theta) &\leq \frac{d}{d\theta} V(\theta, \theta, \chi, \tau) \leq E_{\chi[\theta]} H_{\theta}(p, \theta) - H_{\theta}(p_0, \theta), \\ \Leftrightarrow E_{\chi[\theta]} H_{\theta}(p, \theta) - H_{\theta}(p_0, \theta) &= \frac{d}{d\theta} V(\theta, \theta, \chi, \tau). \end{aligned}$$

(i) follows. Substituting (i) into (9) yields (ii).

(If) Given (i), incentive compatibility (9) reduces to (ii). (i) also implies that $V(\theta, \theta, \chi, \tau)$ is nondecreasing, so the participation constraint (3) reduces to (iii). \square

Proof of Lemma 2 (Reformulated Direct Mechanism Design Problem)

Lemma 1(i) is equivalent to

$$\tau(\theta) = V(0, 0, \chi, \tau) + \int_0^\theta E_{\chi[r]}[H_\theta(p, r) - H_\theta(p_0, r)]dr - E_{\chi[\theta]}[H(p, \theta) - H(p_0, \theta)].$$

Therefore, the objective function in (1) can be equivalently written

$$\int_0^1 E_{\chi[\theta]} [w(p) + H_\theta(p, \theta) - H_\theta(p_0, \theta)] f(\theta) d\theta - \int_0^1 \int_0^\theta E_{\chi[r]} [H_\theta(p, r) - H_\theta(p_0, r)] f(\theta) dr d\theta - V(0, 0, \chi, \tau).$$

Integrating the second term by parts reduces this to

$$\int_0^1 E_{\chi[\theta]} [w(p) + H_\theta(p, \theta) - H_\theta(p_0, \theta)] f(\theta) d\theta - \int_0^1 E_{\chi[s]} [H_\theta(p, s) - H_\theta(p_0, s)] (1 - F(\theta)) d\theta,$$

or equivalently

$$E_F \left(E_{\chi[\theta]} \left(w(p) - \frac{1 - F(\theta)}{f(\theta)} (H_\theta(p, \theta) - H_\theta(p_0, \theta)) + H(p, \theta) - H(p_0, \theta) \right) \right) - V(0, 0, \chi, \tau).$$

Then from Lemma 1, we can write (1) equivalently as

$$\max_x E_F \left[E_{\chi[\theta]} \left[w(p) - \frac{1 - F(\theta)}{f(\theta)} (H_\theta(p, \theta) - H_\theta(p_0, \theta)) + H(p, \theta) - H(p_0, \theta) \right] \right] - V(0, 0, \chi, \tau) \quad (10)$$

$$\text{s.t. } E_{\chi[\theta]} p = p_0 \quad \forall \theta \in [0, 1],$$

$$\int_{\theta'}^\theta E_{\chi[r]} [H_\theta(p, r)] dr \geq \int_{\theta'}^\theta E_{\chi[\theta']} [H_\theta(p, r)] dr \quad \forall \theta, \theta' \in [0, 1],$$

$$V(0, 0, \chi, \tau) \geq 0$$

$$\tau(\theta) = V(0, 0, \chi, \tau) + \int_0^\theta E_{\chi[r]} [H_\theta(p, r) - H_\theta(p_0, r)] dr - E_{\chi[\theta]} [H(p, \theta) - H(p_0, \theta)].$$

Since the objective is decreasing in $V(0, 0, \chi, \tau)$, and the only constraint on it is the low type participation constraint, any solution to (10) must set $V(0, 0, \chi, \tau) = 0$. So (1), (10), and (4) are equivalent, as desired. \square

Proof of Lemma 4 (Tangent Lines and Exploitation of Differentiability) (i),(ii): $\Pi(\cdot, \theta)$ is concave and $\Pi(p, \theta) \geq \pi(p, \theta) \forall p$ by definition. Then from these facts and Jensen's inequality, for any Bayes-plausible x ,

$$\Pi(p_0, \theta) \geq \int \Pi(p, \theta) d\nu(p) \geq \int \pi(p, \theta) d\nu(p).$$

From Lemma 3, ν^* is optimal if and only if both inequalities bind. The first does iff $\Pi(\cdot, \theta)$ is affine on the convex hull of $\text{supp } \nu^*$, since it is concave. Then unless $\text{supp } \nu^* = \{p_0\}$, $\Pi_p(p_0, \theta)$ exists. (i) follows. The second inequality binds iff $\Pi(p, \theta) = \pi(p, \theta)$ ν^* -a.e. Since $\pi(\cdot, \theta)$ is upper semicontinuous, this is equivalent to $\Pi(p, \theta) = \pi(p, \theta) \forall p \in \text{supp } \nu^*$: Obviously the latter implies the former. To show the reverse is true, note that since $\Pi(\cdot, \theta)$ is concave, it is continuous. Then $\pi(\cdot, \theta) - \Pi(\cdot, \theta)$ is upper semicontinuous. Suppose that $\pi(x_0, \theta) - \Pi(x_0, \theta) < 0$ for some $x_0 \in \text{supp } \nu^*$. Then $\pi(p, \theta) - \Pi(p, \theta) < 0$ for all p in some neighborhood $U \ni x_0$. But since $x_0 \in \text{supp } \nu^*$, $\nu^*(U) > 0$, a contradiction.

(iii): Since $\Pi(\cdot, \theta)$ is concave, for all $p' \in [0, 1]$ we have

$$\begin{aligned}\pi(p', \theta) &\leq \Pi(p', \theta) \leq \Pi(p_0, \theta) + \Pi_p(p_0, \theta)(p' - p_0) \\ &\leq \pi(p, \theta) + \Pi_p(p_0, \theta)(p' - p),\end{aligned}$$

where the second inequality follows from (ii). Then for $0 < \epsilon < \min\{p, 1 - p\}$,

$$\begin{aligned}\pi(p + \epsilon, \theta) &\leq \pi(p, \theta) + \Pi_p(p_0, \theta)\epsilon, \\ \pi(p - \epsilon, \theta) &\leq \pi(p, \theta) - \Pi_p(p_0, \theta)\epsilon, \\ \Rightarrow \frac{\pi(p + \epsilon, \theta) - \pi(p, \theta)}{\epsilon} &\leq \Pi_p(p_0, \theta) \leq \frac{\pi(p, \theta) - \pi(p - \epsilon, \theta)}{\epsilon}.\end{aligned}$$

Then by the squeeze theorem, $\Pi_p(p_0, \theta) = \pi_p(p, \theta)$.

(iv): Suppose $\pi_{pp}(p, \theta) > 0$. Then $\pi(\cdot, \theta)$ is strictly convex in a neighborhood of p , so there exists $\epsilon > 0$ such that $\frac{\pi(p+\epsilon, \theta) + \pi(p-\epsilon, \theta)}{2} > \pi(p, \theta)$. But then $(p, \pi(p, \theta)) \in \text{int}(\text{conv}(\pi(\cdot, \theta)))$, so $\pi(p, \theta) < \Pi(p, \theta)$, contradicting (i). \square

Proof of Theorem 2 (Sufficiency and Uniqueness for Double-Peaked Problems) Let p_1, p_2 be the inflection points of $\pi(\cdot, \theta)$.

(i) ($\pi(\cdot, \theta)$ has at most one interior critical pair): Suppose that $\{\underline{p}, \bar{p}\}$ and $\{\underline{p}', \bar{p}'\}$ are both interior critical pairs of $\pi(\cdot, \theta)$. Then we have $\pi_p(\underline{p}, \theta) = \pi_p(\bar{p}, \theta)$, $\pi_p(\underline{p}', \theta) = \pi_p(\bar{p}', \theta)$ and $\pi_{pp}(\underline{p}, \theta), \pi_{pp}(\underline{p}', \theta), \pi_{pp}(\bar{p}, \theta), \pi_{pp}(\bar{p}', \theta) \leq 0$. Accordingly, $\underline{p}, \underline{p}' \in (0, p_1]$ and $\bar{p}, \bar{p}' \in [p_2, 1)$, intervals on the interior of which $\pi(\cdot, \theta)$ is strictly concave. Suppose without loss that $\underline{p} < \underline{p}'$. Then $\pi_p(\underline{p}, \theta) > \pi_p(\underline{p}', \theta) \Rightarrow \pi_p(\bar{p}, \theta) > \pi_p(\bar{p}', \theta) \Rightarrow \bar{p} < \bar{p}'$. Now due to strict concavity, we have

$$\begin{aligned}\pi(\underline{p}, \theta) + \pi_p(\underline{p}, \theta)(\underline{p}' - \underline{p}) &\geq \pi(\underline{p}', \theta), \\ \pi(\bar{p}', \theta) + \pi_p(\bar{p}', \theta)(\bar{p} - \bar{p}') &\geq \pi(\bar{p}, \theta), \\ \pi_p(\underline{p}, \theta)(\bar{p} - \underline{p}') &> \pi_p(\underline{p}', \theta)(\bar{p} - \underline{p}'), \\ \Rightarrow \pi(\underline{p}, \theta) + \pi_p(\underline{p}, \theta)(\bar{p} - \underline{p}) &> \pi(\underline{p}', \theta) + \pi_p(\underline{p}', \theta)(\bar{p} - \underline{p}') = \pi(\bar{p}', \theta) + \pi_p(\bar{p}', \theta)(\bar{p} - \bar{p}') \geq \pi(\bar{p}, \theta),\end{aligned}$$

a contradiction.

The following lemma is necessary for (ii) and (iii):

Lemma 10. *If $\pi(\cdot, \theta)$ is double-peaked, then for any critical pair $\{\underline{p}, \bar{p}\}$ which is not interior, $\{\underline{p}, \bar{p}\} \not\subseteq \text{supp } \nu^*$ for any $\nu^* \in \arg \max_{\nu} \{E_{\nu} \pi(p, \theta) \text{ s.t. } E_{\nu} p = p_0\}$.*

Proof. Without loss, let $\underline{p} = 0$, and suppose that $\text{supp } \nu^* \supseteq \{0, \bar{p}\}$ for some $\nu^* \in \arg \max_{\nu} \{E_{\nu} \pi(p, \theta) \text{ s.t. } E_{\nu} p = p_0\}$. Then from Lemma 4, $\Pi(0, \theta) = \Pi(p_0, \theta) - \Pi_p(p_0, \theta)p_0$. Since $\Pi(\cdot, \theta)$ is concave, this means that for all $p \in [0, p_0]$, $\Pi(p, \theta) = \Pi(0, \theta) + \Pi_p(p_0, \theta)p$. But since $\lim_{z \rightarrow 0} \pi_p(z, \theta) = \infty$, there exists $\epsilon > 0$ such that $\pi_p(\epsilon, \theta) > \Pi_p(p_0, \theta)$ and $\pi_{pp}(\epsilon, \theta) < 0$. Then

$$\Pi(0, \theta) + \Pi_p(p_0, \theta)\epsilon < \Pi(0, \theta) + \pi_p(\epsilon, \theta)p < \pi(\epsilon, \theta),$$

a contradiction. \square

(ii): From Lemma 10, for any $\nu^* \in \arg \max_{\nu} \{E_{\nu} \pi(p, \theta) \text{ s.t. } E_{\nu} p = p_0\}$, the only critical pairs of $\pi(\cdot, \theta)$ which $\text{supp } \nu^*$ can contain are those which are interior. It follows from Theorem 1 (i) that if no interior critical pairs exist, $\text{supp } \nu^* \cap [0, p_0) = \emptyset$ or $\text{supp } \nu^* \cap (p_0, 1] = \emptyset$. From Bayes-plausibility, $\text{supp } \nu^* = \{p_0\}$.

Lemma 11. *If $\pi(\cdot, \theta)$ is double-peaked and has an interior critical pair $\{\underline{p}, \bar{p}\}$, then $p_0 \notin \text{supp } \nu^*$ for any $\nu^* \in \arg \max_{\nu} \{E_{\nu} \pi(p, \theta) \text{ s.t. } E_{\nu} p = p_0\}$.*

Proof. $\Pi(p_0, \theta) \geq \pi(\underline{p}, \theta) + \pi_p(\underline{p}, \theta)(p_0 - \underline{p})$, since $(p_0, \pi(\underline{p}, \theta) + \pi_p(\underline{p}, \theta)(p_0 - \underline{p}))$ can be written as a convex combination of $(\underline{p}, \pi(\underline{p}, \theta))$ and $(\bar{p}, \pi(\bar{p}, \theta))$. Unless $p_0 \in (p_1, p_2)$, strict concavity tells us that $\pi(\underline{p}, \theta) + \pi_p(\underline{p}, \theta)(p_0 - \underline{p}) = \pi(\bar{p}, \theta) + \pi_p(\bar{p}, \theta)(p_0 - \bar{p}) > \pi(p_0, \theta)$ and so $\Pi(p_0, \theta) > \pi(p_0, \theta)$. Otherwise, note that

$$\begin{aligned} \pi(\underline{p}, \theta) + \pi_p(\underline{p}, \theta)(p_1 - \underline{p}) &> \pi(p_1, \theta), && \text{(by strict concavity on } (0, p_1)) \\ \pi(\bar{p}, \theta) + \pi_p(\bar{p}, \theta)(p_2 - \bar{p}) &> \pi(p_2, \theta), && \text{(by strict concavity on } (p_2, 1)) \\ \frac{p_2 - p_0}{p_2 - p_1} \pi(p_1, \theta) + \frac{p_0 - p_1}{p_2 - p_1} \pi(p_2, \theta) &> \pi(p_0, \theta), && \text{(by strict convexity on } (p_1, p_2)) \\ \Rightarrow \Pi(p_0, \theta) &\geq \pi(\underline{p}, \theta) + \pi_p(\underline{p}, \theta)(p_0 - \underline{p}) > \pi(p_0, \theta). \end{aligned}$$

Since in all cases, $\Pi(p_0, \theta) > \pi(p_0, \theta)$, $\{p_0\} \notin \text{supp } \nu^*$ for all $\nu^* \in \arg \max_{\nu} \{E_{\nu} \pi(p, \theta) \text{ s.t. } E_{\nu} p = p_0\}$ by Lemma 4. \square

(iii): From Lemmas 11 and 10 and Theorem 1 (i), for any $\nu^* \in \arg \max_{\nu} \{E_{\nu} \pi(p, \theta) \text{ s.t. } E_{\nu} p = p_0\}$, $\text{supp } \nu^*$ must be an interior critical pair of $\pi(\cdot, \theta)$. Since there is only one interior critical pair of $\pi(\cdot, \theta)$, this ν^* is unique. \square

Proof of Lemma 6 (Monotone Participation) (i): $\Pi(p_0, \theta) = 0$ for all $\theta \in \Theta_0$, so the totally uninformative experiment ν_0 solves $\max_{\nu} \{E_{\nu} \pi(p, \theta) \text{ s.t. } E_{\nu} p = p_0\}$. From Theorem 2 (i) and (ii), it follows that $\pi(\cdot, \theta)$ must have no interior pair.

(ii): By definition, for all $\theta \in \Theta_+$, $\Pi(p_0, \theta) > \pi(p_0, \theta) = 0$. It follows that $p_0 \notin \nu^*$ for all $\nu^* \in \arg \max_{\nu} \{E_{\nu} \pi(p, \theta) \text{ s.t. } E_{\nu} p = p_0\}$. From Theorem 2 (i) and (ii), it follows that $\pi(\cdot, \theta)$ must have a unique interior pair.

(iii): Let $\nu^* \in \arg \max_{\nu} \{E_{\nu} \pi(p, \theta) \text{ s.t. } E_{\nu} p = p_0\}$. Then we have

$$\begin{aligned} \Pi(p_0, \theta) &= E_{\nu^*} \pi(p, \theta) \\ &\leq E_{\nu^*} \pi(p, \theta') && \text{(by Jensen's inequality)} \\ &\leq \max_{\nu} \{E_{\nu} \pi(p, \theta') \text{ s.t. } E_{\nu} p = p_0\} \\ &= \Pi(p_0, \theta'). \end{aligned}$$

Since $\Pi(p_0, \theta)$ is a continuous function of θ by the maximum theorem, Θ_+ is open; thus, if $\Theta_+ \neq \Theta$ and $\Theta_+ \neq \emptyset$, it follows that $\Theta_+ = (\underline{\theta}, 1]$. \square

Proof of Theorem 3 (Convexity Comparative Statics) I require the following lemma:

Lemma 12 (Change in Critical Pairs). *Suppose that π is smoothly double-peaked. If an interior critical pair $\{\underline{p}(\theta), \bar{p}(\theta)\}$ of $\pi(\cdot, \theta)$ exists for all θ in an open interval about θ_0 , then the functions \underline{p} and \bar{p} are continuously differentiable in a neighborhood of θ_0 , and their derivatives are given by*

$$\begin{aligned}\underline{p}'(\theta) &= \frac{\pi_\theta(\bar{p}(\theta), \theta) - \pi_\theta(\underline{p}(\theta), \theta) - \pi_{p\theta}(\underline{p}(\theta), \theta)(\bar{p}(\theta) - \underline{p}(\theta))}{\pi_{pp}(\underline{p}(\theta), \theta)(\bar{p}(\theta) - \underline{p}(\theta))}, \\ \bar{p}'(\theta) &= \frac{\pi_\theta(\bar{p}(\theta), \theta) - \pi_\theta(\underline{p}(\theta), \theta) - \pi_{p\theta}(\bar{p}(\theta), \theta)(\bar{p}(\theta) - \underline{p}(\theta))}{\pi_{pp}(\bar{p}(\theta), \theta)(\bar{p}(\theta) - \underline{p}(\theta))}.\end{aligned}$$

Proof. First, no critical pair of a double-peaked value function can have an inflection point of the value function as one of its elements: Suppose without loss that $p_1(\theta), \bar{p}(\theta)$ is a critical pair of $\pi(\cdot, \theta)$. Since $\pi(\cdot, \theta)$ is strictly convex on $(p_1(\theta), p_2(\theta))$, the tangent line to $\pi(\cdot, \theta)$ at $p_1(\theta)$ must lie strictly below the value function on $(p_1(\theta), p_2(\theta))$. Since the tangent line must intersect $\pi(\cdot, \theta)$ at $\bar{p}(\theta)$, we must have $\bar{p}(\theta) > p_2(\theta)$. But $\pi(\cdot, \theta)$ is strictly concave at such $\bar{p}(\theta)$, and so the tangent line to $\pi(\cdot, \theta)$ at $\bar{p}(\theta)$ must lie strictly above the value function at $p_2(\theta)$. Then the two tangent lines cannot be the same, a contradiction.

Thus, we must have $\pi_{pp}(\underline{p}(\theta), \theta) < 0$ and $\pi_{pp}(\bar{p}(\theta), \theta) < 0$, and the derivatives follow immediately from the definition of critical pair, Theorem 2, and the implicit function theorem. \square

For (i), note that since $\pi_{pp}(\bar{p}(\theta), \theta)(\bar{p}(\theta) - \underline{p}(\theta)) < 0$,

$$\begin{aligned}\bar{p}'(\theta) \geq 0 &\Leftrightarrow \pi_\theta(\bar{p}(\theta), \theta) - \pi_\theta(\underline{p}(\theta), \theta) \leq \pi_{p\theta}(\bar{p}(\theta), \theta)(\bar{p}(\theta) - \underline{p}(\theta)), \\ &\Leftrightarrow \int_{\underline{p}(\theta)}^{\bar{p}(\theta)} \pi_{\theta p}(z, \theta) dz \leq \int_{\underline{p}(\theta)}^{\bar{p}(\theta)} \pi_{p\theta}(\bar{p}(\theta), \theta) dz,\end{aligned}$$

which follows from $\pi_{\theta p}$ nondecreasing in p ; $\underline{p}'(\theta) \leq 0$ by identical reasoning. For (ii), note that if $\pi_{\theta p}$ is increasing in p , the final inequality is strict and $\bar{p}'(\theta) > 0$; likewise, $\underline{p}'(\theta) < 0$. \square

Proof of Lemma 7 If $\Theta_+ = \Theta$ then $\lim_{\theta \rightarrow \underline{\theta}} \underline{p}(\theta) - \bar{p}(\theta) = \underline{p}(0) - \bar{p}(0) > 0$ and we are done. Otherwise, note that \underline{p}, \bar{p} are continuously differentiable on $(\underline{\theta}, 1)$ and so have a Lipschitz extension to $[\underline{\theta}, 1]$; thus, the limit exists.

For each $(p, \theta) \in [0, 1] \times \Theta$ with $p < p_1(\theta)$, let $Z_1(p, \theta)$ be some neighborhood of (p, θ) on which $\pi_{pp} < 0$ (which exists, since each such (p, θ) has a neighborhood on which π_{pp} is continuous). Likewise, for each (p, θ) with $p_1(\theta) < p < p_2(\theta)$, let $Z_+(p, \theta)$ be some neighborhood of (p, θ) on which $\pi_{pp} > 0$; and for each (p, θ) with $p > p_2(\theta)$, let $Z_2(p, \theta)$ be some neighborhood of (p, θ) on which $\pi_{pp} < 0$.

Further, let $Y_1 = \bigcup_{(p, \theta): p < p_1(\theta)} Z_1(p, \theta)$, $Y_2 = \bigcup_{(p, \theta): p > p_2(\theta)} Z_2(p, \theta)$, $Y_+ = \bigcup_{(p, \theta): p_1(\theta) < p < p_2(\theta)} Z_+(p, \theta)$. Y_1 , Y_2 , and Y_+ are each open and disjoint from one another. Then $Y_+ \cap \text{cl } Y_1 = Y_2 \cap \text{cl } Y_1 = \emptyset$ and $Y_+ \cap \text{cl } Y_2 = Y_1 \cap \text{cl } Y_2 = \emptyset$. Note the following facts about the graphs of p_1 and p_2 :

- $\text{Gr } p_1 \subset \text{cl } Y_1$, since every neighborhood of $(\theta, p_1(\theta))$ contains (θ, p) with $p < p_1(\theta)$. Likewise, $\text{Gr } p_2 \subset \text{cl } Y_2$, $\text{Gr } p_1 \subset \text{cl } Y_+$, and $\text{Gr } p_2 \subset \text{cl } Y_+$.

- $\text{Gr } p_1 \cap Y_1 = \emptyset$, since every point in $\text{Gr } p_1$ is a limit point of Y_+ , which is disjoint from Y_1 . Likewise, $\text{Gr } p_1 \cap Y_+ = \text{Gr } p_1 \cap Y_2 = \emptyset$ and $\text{Gr } p_2 \cap Y_2 = \text{Gr } p_2 \cap Y_+ = \text{Gr } p_2 \cap Y_1 = \emptyset$.
- $Y_1 \cup \text{Gr } p_1 \cup Y_+ \cup \text{Gr } p_2 \cup Y_2 = \Theta \times [0, 1]$.
- $\text{Gr } p_2 \cap \text{cl } Y_1 = \emptyset$: Suppose $(p_2(\theta'), \theta')$ is a limit point of Y_1 . Choose $\epsilon \in (0, p_2(\theta') - p_1(\theta'))$. Then $(p_2(\theta') - \epsilon, \theta') \in Y_+$. Since Y_+ is open, there exists $\delta > 0$ such that the δ -ball about $(p_2(\theta') - \epsilon, \theta')$ is a subset of Y_+ . Then for all $\delta' < \delta$, $p_2(\theta') - \epsilon > p_1(\theta' - \delta')$ and $p_2(\theta') - \epsilon > p_1(\theta' + \delta')$. Then $(p_2(\theta') - \epsilon, 1] \times (\theta' - \delta, \theta' + \delta)$ does not intersect Y_1 , a contradiction.
- Likewise, $\text{Gr } p_1 \cap \text{cl } Y_2 = \emptyset$.

It follows that $\text{Gr } p_1 \cup Y_1 = \text{cl } Y_1$. Then its complement in $[0, 1] \times \Theta$, $Y_+ \cup \text{Gr } p_2 \cup Y_2$, is open. Then $\text{Gr } p_1$ is closed, and so p_1 is continuous, by the closed graph theorem; likewise, p_2 is continuous. Then $\lim_{\theta \rightarrow \underline{\theta}} \bar{p}(\theta) - \underline{p}(\theta) > \lim_{\theta \rightarrow \underline{\theta}} p_2(\theta) - p_1(\theta) = p_2(\underline{\theta}) - p_1(\underline{\theta}) > 0$, as desired. \square

Proof of Theorem 5 (Sufficient Condition for a Global Maximum) Suppose that

$\nu^* \notin \arg \max_{\nu} \{E_{\nu} v(p, \theta, \psi) \text{ s.t. } E_{\nu} p = p_0\}$. By Theorem 1, there exists some other critical pair $\{\underline{p}', \bar{p}'\}$ which supports a solution. Then

$$\begin{aligned}
\max_{\nu} \{E_{\nu} v(p, \theta, \psi) \text{ s.t. } E_{\nu} p = p_0\} &= \frac{\bar{p}' - p_0}{\bar{p}' - \underline{p}'} v(\underline{p}', \theta, \psi) + \frac{p_0 - \underline{p}'}{\bar{p}' - \underline{p}'} v(\bar{p}', \theta, \psi) \\
&\leq v(\underline{p}, \theta, \psi) + \left(\frac{\bar{p}' - p_0}{\bar{p}' - \underline{p}'} (\underline{p}' - \underline{p}) + \frac{p_0 - \underline{p}'}{\bar{p}' - \underline{p}'} (\bar{p}' - \underline{p}) \right) \frac{v(y, \theta, \psi) - v(\underline{p}, \theta, \psi)}{\bar{p} - \underline{p}} \\
&= v(\underline{p}, \theta, \psi) + \frac{p_0 - \underline{p}}{\bar{p} - \underline{p}} (v(\bar{p}, \theta, \psi) - v(\underline{p}, \theta, \psi)) \\
&= \frac{\bar{p} - p_0}{\bar{p} - \underline{p}} v(\underline{p}, \theta, \psi) + \frac{p_0 - \underline{p}}{\bar{p} - \underline{p}} v(\bar{p}, \theta, \psi) \\
&= E_{\nu^*} v(p, \theta, \psi).
\end{aligned}$$

a contradiction. \square

Proof of Lemma 8 (Sufficient Conditions for Publication Implementation) For $\theta > \tilde{\theta}$, implementation of χ follows from (5)-(7) and Theorems 1 and 5.

If $\text{supp } \chi[\tilde{\theta}] = \{\bar{y}(\tilde{\theta}), \underline{y}(\tilde{\theta})\}$, the same follows for $\tilde{\theta}$.

If $\text{supp } \chi[\tilde{\theta}] = \{p_0\}$, then $\sup\{z | (p_0, z) \in \text{conv}(v(\cdot, \tilde{\theta}, \psi))\} = 0$:

$$\begin{aligned}
&(p_0 - \underline{y}(\tilde{\theta})) \left(\psi(\bar{y}(\tilde{\theta})) + H(\bar{y}(\tilde{\theta}), \tilde{\theta}) - H(p_0, \tilde{\theta}) \right) + (\bar{y}(\tilde{\theta}) - p_0) \left(\psi(\underline{y}(\tilde{\theta})) + H(\underline{y}(\tilde{\theta}), \tilde{\theta}) - H(p_0, \tilde{\theta}) \right) = 0 \tag{8} \\
\Leftrightarrow (p_0 - \underline{y}(\tilde{\theta})) \left((H_p(\underline{y}(\tilde{\theta}), \tilde{\theta}) + \psi'(\underline{y}(\tilde{\theta}))) (\bar{y}(\tilde{\theta}) - \underline{y}(\tilde{\theta})) \right) + (\bar{y}(\tilde{\theta}) - \underline{y}(\tilde{\theta})) \left(\psi(\underline{y}(\tilde{\theta})) + H(\underline{y}(\tilde{\theta}), \tilde{\theta}) - H(p_0, \tilde{\theta}) \right) = 0 \tag{from (7)} \\
\Leftrightarrow (\bar{y}(\tilde{\theta}) - \underline{y}(\tilde{\theta})) \left((p_0 - \underline{y}(\tilde{\theta})) (H_p(\underline{y}(\tilde{\theta}), \tilde{\theta}) + \psi'(\underline{y}(\tilde{\theta}))) + \psi(\underline{y}(\tilde{\theta})) + H(\underline{y}(\tilde{\theta}), \tilde{\theta}) - H(p_0, \tilde{\theta}) \right) = 0 \\
\Rightarrow (p_0 - p) (H_p(\underline{y}(\theta), \theta) + \psi'(\underline{y}(\theta))) + \psi(p) + H(p, \tilde{\theta}) - H(p_0, \tilde{\theta}) \leq 0 \quad \forall p \in [0, 1] \text{ from (5)}.
\end{aligned}$$

Then for all $\underline{p} < p_0 < \bar{p}$,

$$\begin{aligned} & \frac{\bar{p} - p_0}{\bar{p} - \underline{p}} (\psi(\underline{p}) + H(\underline{p}, \tilde{\theta}) - H(p_0, \tilde{\theta})) + \frac{p_0 - \underline{p}}{\bar{p} - \underline{p}} (\psi(\bar{p}) + H(\bar{p}, \tilde{\theta}) - H(p_0, \tilde{\theta})) \\ & \leq \left(\frac{\bar{p} - p_0}{\bar{p} - \underline{p}} (\underline{p} - p_0) + \frac{p_0 - \underline{p}}{\bar{p} - \underline{p}} (\bar{p} - p_0) \right) (H_p(\underline{y}(\theta), \theta) + \psi'(\underline{y}(\theta))) \leq 0. \end{aligned}$$

Thus, $\max_{\nu} \{E_{\nu}[v(p, \tilde{\theta}, \psi)] \text{ s.t. } E_{\nu} p = p_0\} = 0$, implying that the totally uninformative experiment solves type $\tilde{\theta}$'s problem.

For $\theta < \tilde{\theta}$, note that $E_{\chi[\tilde{\theta}]}[H(p, \tilde{\theta}) - H(p_0, \tilde{\theta}) + \psi(p)] = 0$. Since $H_{pp\theta}(p, \theta) > 0$ for $p \neq p_0$, any $\theta < \tilde{\theta}$ receives a strictly worse expected payoff than $\tilde{\theta}$ from performing any informative experiment, and so declines to do so. The expected cost assertion then follows from revenue equivalence (Corollary 1). \square

Proof of Theorem 6 (Implementation Using a Results-Based Contract) The proof shows that ψ_{χ} is differentiable on $[0, \underline{y}(\tilde{\theta})]$ and $(\bar{y}(\tilde{\theta}), 1]$, and for all $\theta \geq \tilde{\theta}$ satisfies conditions (5)-(8); the proposition then follows from Lemma 8.

Step 1: Calculation of $\psi'_{\chi}(p)$ and $\psi''_{\chi}(p)$. Since the restrictions of \underline{y} and \bar{y} to $[\tilde{\theta}, 1]$ are injective, their inverses $\underline{y}^{-1} : [\underline{y}(1), \underline{y}(\tilde{\theta})] \rightarrow [\tilde{\theta}, 1]$ and $\bar{y}^{-1} : [\bar{y}(\tilde{\theta}), \bar{y}(1)] \rightarrow [\tilde{\theta}, 1]$ exist. By the inverse function theorem, since \underline{y} and \bar{y} are continuously differentiable with nonzero derivative on $(\tilde{\theta}, 1)$, the derivatives of these inverses are given by $d\underline{y}^{-1}(p)/dp = (\underline{y}'(\underline{y}^{-1}(p)))^{-1}$ on $(\underline{y}(1), \underline{y}(\tilde{\theta}))$ and $d\bar{y}^{-1}(p)/dp = (\bar{y}'(\bar{y}^{-1}(p)))^{-1}$ on $(\bar{y}(\tilde{\theta}), \bar{y}(1))$.

Applying Leibniz' rule, we have

$$\psi'_{\chi}(p) = \begin{cases} H_p(p_0, \tilde{\theta}) - H_p(\bar{y}(1), 1) + \int_{\tilde{\theta}}^1 \frac{H_{\theta}(\bar{y}(t), t) - H_{\theta}(\underline{y}(t), t)}{\bar{y}(t) - \underline{y}(t)} dt, & p \in [\bar{y}(1), 1), \\ H_p(p_0, \tilde{\theta}) - H_p(p, \bar{y}^{-1}(p)) + \int_{\tilde{\theta}}^{\bar{y}^{-1}(p)} \frac{H_{\theta}(\bar{y}(t), t) - H_{\theta}(\underline{y}(t), t)}{\bar{y}(t) - \underline{y}(t)} dt, & p \in (\bar{y}(\tilde{\theta}), \bar{y}(1)), \\ 0, & p \in (\underline{y}(\tilde{\theta}), \bar{y}(\tilde{\theta})), \\ H_p(p_0, \tilde{\theta}) - H_p(p, \underline{y}^{-1}(p)) + \int_{\tilde{\theta}}^{\underline{y}^{-1}(p)} \frac{H_{\theta}(\bar{y}(t), t) - H_{\theta}(\underline{y}(t), t)}{\bar{y}(t) - \underline{y}(t)} dt, & p \in (\underline{y}(1), \underline{y}(\tilde{\theta})), \\ H_p(p_0, \tilde{\theta}) - H_p(\underline{y}(1), 1) + \int_{\tilde{\theta}}^1 \frac{H_{\theta}(\bar{y}(t), t) - H_{\theta}(\underline{y}(t), t)}{\bar{y}(t) - \underline{y}(t)} dt, & p \in (0, \underline{y}(1)]. \end{cases}$$

$$\psi''_{\chi}(p) = \begin{cases} 0, & p \in (\bar{y}(1), 1), \\ -H_{pp}(p, \bar{y}^{-1}(p)) + \frac{H_{\theta}(p, \bar{y}^{-1}(p)) - H_{\theta}(\underline{y}(\bar{y}^{-1}(p)), \bar{y}^{-1}(p)) - H_{\theta p}(p, \bar{y}^{-1}(p))(p - \underline{y}(\bar{y}^{-1}(p)))}{(p - \underline{y}(\bar{y}^{-1}(p)))\bar{y}'(\bar{y}^{-1}(p))}, & p \in (\bar{y}(\tilde{\theta}), \bar{y}(1)), \\ 0, & p \in (\underline{y}(\tilde{\theta}), \bar{y}(\tilde{\theta})), \\ -H_{pp}(p, \underline{y}^{-1}(p)) + \frac{H_{\theta}(p, \underline{y}^{-1}(p)) - H_{\theta}(\bar{y}(\underline{y}^{-1}(p)), \underline{y}^{-1}(p)) - H_{\theta p}(p, \underline{y}^{-1}(p))(\bar{y}(\underline{y}^{-1}(p)) - p)}{(\bar{y}(\underline{y}^{-1}(p)) - p)\underline{y}'(\underline{y}^{-1}(p))}, & p \in (\underline{y}(1), \underline{y}(\tilde{\theta})), \\ 0, & p \in (0, \underline{y}(1)). \end{cases}$$

Step 2: For all $\theta \in [\tilde{\theta}, 1]$, $(\underline{p}, \bar{p}) = (\underline{y}(\theta), \bar{y}(\theta))$ solves

$$\begin{aligned} H(\underline{p}, \theta) + \psi_{\chi}(\underline{p}) + (H_p(\bar{p}, \theta) + \psi'_{\chi}(\bar{p}))(\bar{p} - \underline{p}) &= H(\bar{p}, \theta) + \psi_{\chi}(\bar{p}), \\ H(\underline{p}, \theta) + \psi_{\chi}(\underline{p}) + (H_p(\underline{p}, \theta) + \psi'_{\chi}(\underline{p}))(\bar{p} - \underline{p}) &= H(\bar{p}, \theta) + \psi_{\chi}(\bar{p}). \end{aligned}$$

Since $H_p(\bar{y}(\theta), \theta) + \psi'_{\chi}(\bar{y}(\theta)) = H_p(p_0, \tilde{\theta}) + \int_{\tilde{\theta}}^{\theta} \frac{H_{\theta}(\bar{y}(t), t) - H_{\theta}(\underline{y}(t), t)}{\bar{y}(t) - \underline{y}(t)} dt = H_p(\underline{y}(\theta), \theta) + \psi'_{\chi}(\underline{y}(\theta))$ for all $\theta \in [\tilde{\theta}, 1]$,

we need only consider the first equation in the system. We have

$$\begin{aligned}
& H(\underline{y}(\theta), \theta) + \psi_\chi(\underline{y}(\theta)) + (H_p(\bar{y}(\theta), \theta) + \psi'_\chi(\bar{y}(\theta)))(\bar{y}(\theta) - \underline{y}(\theta)) - H(\bar{y}, \theta) - \psi_\chi(\bar{y}) = \\
& \int_{\tilde{\theta}}^{\theta} H_p(\bar{y}(t), t)\bar{y}'(t) - H_p(\underline{y}(t), t)\underline{y}'(t) + \frac{H_\theta(\bar{y}(t), t) - H_\theta(\underline{y}(t), t)}{\bar{y}(t) - \underline{y}(t)}(\underline{y}(\theta) - \bar{y}(\theta) + \bar{y}(t) - \underline{y}(t))dt \\
& + H(\bar{y}(\tilde{\theta}), \tilde{\theta}) - H(\underline{y}(\tilde{\theta}), \tilde{\theta}) + H(\underline{y}(\theta), \theta) - H(\bar{y}(\theta), \theta) + H_p(p_0, \tilde{\theta})(\underline{y}(\theta) - \bar{y}(\theta)) \\
& + H_p(p_0, \tilde{\theta})(\bar{y}(\theta) - \underline{y}(\theta)) + \int_{\tilde{\theta}}^{\theta} \frac{H_\theta(\bar{y}(t), t) - H_\theta(\underline{y}(t), t)}{\bar{y}(t) - \underline{y}(t)}(\bar{y}(\theta) - \underline{y}(\theta))dt \\
& = \int_{\tilde{\theta}}^{\theta} H_p(\bar{y}(t), t)\bar{y}'(t) - H_p(\underline{y}(t), t)\underline{y}'(t) + H_\theta(\bar{y}(t), t) - H_\theta(\underline{y}(t), t)dt \\
& + H(\bar{y}(\tilde{\theta}), \tilde{\theta}) - H(\underline{y}(\tilde{\theta}), \tilde{\theta}) + H(\underline{y}(\theta), \theta) - H(\bar{y}(\theta), \theta) \\
& = 0, \text{ (by the multivariate chain rule)}
\end{aligned}$$

as desired. Thus, ψ_χ satisfies (6) and (7).

Step 3: For all $\theta \geq \tilde{\theta}$, $\psi''_\chi(p) + H_{pp}(p, \theta) < 0$ for $p \in [0, \underline{y}(1)] \cup (\underline{y}(1), \underline{y}(\theta)) \cup (\bar{y}(\theta), \bar{y}(1)) \cup (\bar{y}(1), 1]$. Hence, $\psi_\chi(p) + H(p, \theta)$ is concave in p on $p \in [0, \underline{y}(\theta)]$ and on $[\bar{y}(\theta), 1]$.

Since $H_{pp} < 0$ and $\psi''_\chi(p) = 0$ for $p \in [0, \underline{y}(1)] \cup (\bar{y}(1), 1]$, $\psi''_\chi(p) + H_{pp}(p, \theta) < 0$ on those intervals. For all $\theta' > \tilde{\theta}$, we have

$$\begin{aligned}
\psi''_\chi(\bar{y}(\theta')) + H_{pp}(\bar{y}(\theta'), \theta') &= \frac{H_\theta(\bar{y}(\theta'), \theta') - H_\theta(\underline{y}(\theta'), \theta') - H_{\theta p}(\bar{y}(\theta'), \theta')(\bar{y}(\theta') - \underline{y}(\theta'))}{(\bar{y}(\theta') - \underline{y}(\theta'))\bar{y}'(\theta')} \\
&= \frac{\int_{\underline{y}(\theta')}^{\bar{y}(\theta')} H_{\theta p}(p, \theta') - H_{\theta p}(\bar{y}(\theta'), \theta')dp}{(\bar{y}(\theta') - \underline{y}(\theta'))\bar{y}'(\theta')} < 0, \\
\psi''_\chi(\underline{y}(\theta')) + H_{pp}(\underline{y}(\theta'), \theta') &= \frac{H_\theta(\bar{y}(\theta'), \theta') - H_\theta(\underline{y}(\theta'), \theta') - H_{\theta p}(\underline{y}(\theta'), \theta')(\bar{y}(\theta') - \underline{y}(\theta'))}{(\bar{y}(\theta') - \underline{y}(\theta'))\underline{y}'(\theta')} \\
&= \frac{\int_{\underline{y}(\theta')}^{\bar{y}(\theta')} H_{\theta p}(p, \theta') - H_{\theta p}(\underline{y}(\theta'), \theta')dp}{(\bar{y}(\theta') - \underline{y}(\theta'))\underline{y}'(\theta')} < 0.
\end{aligned}$$

Since $H_{pp}(p, \theta) \leq H_{pp}(p, \theta')$ for all $\theta \leq \theta'$, it follows that $\psi''_\chi(\bar{y}(\theta')) + H_{pp}(\bar{y}(\theta'), \theta) < 0$ and $\psi''_\chi(\underline{y}(\theta')) + H_{pp}(\underline{y}(\theta'), \theta) < 0$ for all $\theta \leq \theta'$; equivalently, $\psi''_\chi(p) + H_{pp}(p, \theta) < 0$ for $p \in (\underline{y}(1), \underline{y}(\theta)] \cup [\bar{y}(\theta), \bar{y}(1)]$ for all $\theta \geq \tilde{\theta}$.

Step 4: For all $\theta \geq \tilde{\theta}$,

$$\begin{aligned}
H(\underline{y}(\theta), \theta) + \psi_\chi(\underline{y}(\theta)) + (H_p(\underline{y}(\theta), \theta) + \psi'_\chi(\underline{y}(\theta)))(p - \underline{y}(\theta)) &\geq H(p, \theta) + \psi_\chi(p) \quad \forall p \in [\underline{y}(\theta), \underline{y}(\tilde{\theta})], \\
H(\bar{y}(\theta), \theta) + \psi_\chi(\bar{y}(\theta)) + (H_p(\bar{y}(\theta), \theta) + \psi'_\chi(\bar{y}(\theta)))(p - \bar{y}(\theta)) &\geq H(p, \theta) + \psi_\chi(p) \quad \forall p \in [\bar{y}(\tilde{\theta}), \bar{y}(\theta)].
\end{aligned}$$

Consider the second of these inequalities, which (since $p = \bar{y}(\theta')$ for some $\theta' < \theta$) can readily be seen to be

equivalent to

$$\begin{aligned}
& H(\bar{y}(\theta), \theta) + H_p(p_0, \tilde{\theta})\bar{y}'(\theta') + \int_{\tilde{\theta}}^{\theta} \frac{H_{\theta}(\bar{y}(t), t) - H_{\theta}(\underline{y}(t), t)}{\bar{y}(t) - \underline{y}(t)} (\bar{y}(\theta') - \bar{y}(t)) - H_p(\bar{y}(t), t)\bar{y}'(t) dt \\
& \geq H(\bar{y}(\theta'), \theta) + \bar{y}(\theta')H_p(p_0, \tilde{\theta}) + \int_{\tilde{\theta}}^{\theta'} \frac{H_{\theta}(\bar{y}(t), t) - H_{\theta}(\underline{y}(t), t)}{\bar{y}(t) - \underline{y}(t)} (\bar{y}(\theta') - \bar{y}(t)) - H_p(\bar{y}(t), t)\bar{y}'(t) dt \\
\Leftrightarrow & H(\bar{y}(\theta), \theta) - H(\bar{y}(\theta'), \theta) + \int_{\theta'}^{\theta} \frac{H_{\theta}(\bar{y}(t), t) - H_{\theta}(\underline{y}(t), t)}{\bar{y}(t) - \underline{y}(t)} (\bar{y}(\theta') - \bar{y}(t)) - H_p(\bar{y}(t), t)\bar{y}'(t) dt \geq 0 \\
\Leftrightarrow & \int_{\theta'}^{\theta} \frac{H_{\theta}(\bar{y}(t), t) - H_{\theta}(\underline{y}(t), t)}{\bar{y}(t) - \underline{y}(t)} (\bar{y}(\theta') - \bar{y}(t)) + H_p(\bar{y}(t), \theta)\bar{y}'(t) - H_p(\bar{y}(t), t)\bar{y}'(t) dt \geq 0 \\
\Leftrightarrow & \int_{\theta'}^{\theta} \frac{H_{\theta}(\bar{y}(t), t) - H_{\theta}(\underline{y}(t), t)}{\bar{y}(t) - \underline{y}(t)} (\bar{y}(\theta') - \bar{y}(t)) + \int_t^{\theta} H_{p\theta}(\bar{y}(t), r)\bar{y}'(t) dr dt \geq 0.
\end{aligned}$$

Applying multiplicative separability, we can write this as

$$\int_{\theta'}^{\theta} (-g'(t)) \frac{h(\bar{y}(t)) - h(\underline{y}(t))}{\bar{y}(t) - \underline{y}(t)} (\bar{y}(t) - \bar{y}(\theta')) dt + \int_{\theta'}^{\theta} h'(\bar{y}(t))\bar{y}'(t) \int_t^{\theta} g'(r) dr dt \geq 0.$$

Integrating the second integral by parts yields

$$\begin{aligned}
& \int_{\theta'}^{\theta} (-g'(t)) \frac{h(\bar{y}(t)) - h(\underline{y}(t))}{\bar{y}(t) - \underline{y}(t)} (\bar{y}(t) - \bar{y}(\theta')) dt + \int_{\theta'}^{\theta} g'(t)h(\bar{y}(t)) - g'(t)h(\bar{y}(\theta')) dt \geq 0 \\
\Leftrightarrow & \int_{\theta'}^{\theta} \frac{(-g'(t))}{\bar{y}(t) - \underline{y}(t)} \left((\bar{y}(t) - \bar{y}(\theta')) \int_{\underline{y}(t)}^{\bar{y}(t)} h'(z) dz - (\bar{y}(t) - \underline{y}(t)) \int_{\bar{y}(\theta')}^{\bar{y}(t)} h'(z) dz \right) \geq 0 \\
\Leftrightarrow & \int_{\theta'}^{\theta} \frac{(-g'(t))}{\bar{y}(t) - \underline{y}(t)} \left((\bar{y}(t) - \bar{y}(\theta')) \int_{\underline{y}(t)}^{\bar{y}(\theta')} h'(z) dz - (\bar{y}(\theta') - \underline{y}(t)) \int_{\bar{y}(\theta')}^{\bar{y}(t)} h'(z) dz \right) \geq 0.
\end{aligned}$$

$h''(p) < 0$, so $h'(z) \geq h'(\bar{y}(\theta'))$ for $z \leq \bar{y}(\theta')$ (as in the first inner integral) and $-h'(z) \geq -h'(\bar{y}(\theta'))$ for $z \geq \bar{y}(\theta')$ (as in the second inner integral). Since $g'(t) < 0$,

$$\begin{aligned}
& \int_{\theta'}^{\theta} \frac{(-g'(t))}{\bar{y}(t) - \underline{y}(t)} \left((\bar{y}(t) - \bar{y}(\theta')) \int_{\underline{y}(t)}^{\bar{y}(\theta')} h'(z) dz - (\bar{y}(\theta') - \underline{y}(t)) \int_{\bar{y}(\theta')}^{\bar{y}(t)} h'(z) dz \right) \\
& \geq \int_{\theta'}^{\theta} \frac{(-g'(t))}{\bar{y}(t) - \underline{y}(t)} \left((\bar{y}(t) - \bar{y}(\theta')) \int_{\underline{y}(t)}^{\bar{y}(\theta')} h'(\bar{y}(\theta')) dz - (\bar{y}(\theta') - \underline{y}(t)) \int_{\bar{y}(\theta')}^{\bar{y}(t)} h'(\bar{y}(\theta')) dz \right) \\
& = 0.
\end{aligned}$$

The first inequality follows similarly.

Step 5: For all $\theta \geq \tilde{\theta}$,

$$H(\bar{y}(\theta), \theta) + \psi_{\chi}(\bar{y}(\theta)) + (H_p(\bar{y}(\theta), \theta) + \psi'_{\chi}(\bar{y}(\theta)))(p - \bar{y}(\theta)) \geq H(p, \theta) \quad \forall p \in [\underline{y}(\tilde{\theta}), \bar{y}(\tilde{\theta})].$$

The above inequality is equivalent to

$$\begin{aligned}
& H(\bar{y}(\theta), \theta) - H(p, \theta) + H(p_0, \tilde{\theta}) - H(\bar{y}(\tilde{\theta}), \tilde{\theta}) + (\bar{y}(\theta) - p_0)H_p(p_0, \tilde{\theta}) \\
& \quad + \int_{\tilde{\theta}}^{\theta} \frac{H_{\theta}(\bar{y}(t), t) - H_{\theta}(\underline{y}(t), t)}{\bar{y}(t) - \underline{y}(t)} (p - \bar{y}(t)) - H_p(\bar{y}(t), t)\bar{y}'(t) dt \\
& \quad + H_p(p_0, \tilde{\theta})(p - \bar{y}(\theta)) \geq 0 \\
& \Leftrightarrow (p - p_0)H_p(p_0, \tilde{\theta}) - \int_{p_0}^p H_p(z, \tilde{\theta}) dz - \int_{\tilde{\theta}}^{\theta} H_{\theta}(p, t) dt + \int_{\tilde{\theta}}^{\theta} H_p(\bar{y}(t), t)\bar{y}'(t) + H_{\theta}(\bar{y}(t), t) dt \\
& \quad + \int_{\tilde{\theta}}^{\theta} \frac{H_{\theta}(\bar{y}(t), t) - H_{\theta}(\underline{y}(t), t)}{\bar{y}(t) - \underline{y}(t)} (p - \bar{y}(t)) - H_p(\bar{y}(t), t)\bar{y}'(t) dt \geq 0 \\
& \Leftrightarrow - \int_{p_0}^p \int_{p_0}^z H_{pp}(r, \tilde{\theta}) dr dz + \int_{\tilde{\theta}}^{\theta} \frac{\int_{\underline{y}(t)}^{\bar{y}(t)} H_{\theta p}(z, t) dz (p - \bar{y}(t)) + \int_p^{\bar{y}(t)} H_{\theta p}(z, t) dz (\bar{y}(t) - \underline{y}(t))}{\bar{y}(t) - \underline{y}(t)} dt \geq 0 \\
& \Leftrightarrow - \int_{p_0}^p \int_{p_0}^z H_{pp}(r, \tilde{\theta}) dr dz + \int_{\tilde{\theta}}^{\theta} \frac{- \int_{\underline{y}(t)}^p H_{\theta p}(z, t) dz (\bar{y}(t) - p) + \int_p^{\bar{y}(t)} H_{\theta p}(z, t) dz (p - \underline{y}(t))}{\bar{y}(t) - \underline{y}(t)} dt \geq 0.
\end{aligned}$$

Since $H_{pp\theta} > 0$, $-H_{\theta p}(z, t) \geq -H_{\theta p}(p, t)$ for $z \leq p$ (as in the first inner integral) and $H_{\theta p}(z, t) \geq H_{\theta p}(p, t)$ for $z \geq p$ (as in the second inner integral). From this and $H_{pp} < 0$,

$$\begin{aligned}
& - \int_{p_0}^p \int_{p_0}^z H_{pp}(r, \tilde{\theta}) dr dz + \int_{\tilde{\theta}}^{\theta} \frac{- \int_{\underline{y}(t)}^p H_{\theta p}(z, t) dz (\bar{y}(t) - p) + \int_p^{\bar{y}(t)} H_{\theta p}(z, t) dz (p - \underline{y}(t))}{\bar{y}(t) - \underline{y}(t)} dt \\
& \geq \int_{\tilde{\theta}}^{\theta} \frac{- \int_{\underline{y}(t)}^p H_{\theta p}(p, t) dz (\bar{y}(t) - p) + \int_p^{\bar{y}(t)} H_{\theta p}(p, t) dz (p - \underline{y}(t))}{\bar{y}(t) - \underline{y}(t)} dt \\
& = 0,
\end{aligned}$$

as desired.

Steps 2 through 5 then combine to show that

$$H(\underline{y}(\theta), \theta) + \psi_{\chi}(\underline{y}(\theta)) + (H_p(\underline{y}(\theta), \theta) + \psi'_{\chi}(\underline{y}(\theta)))(p - \underline{y}) \geq H(p, \theta) + \psi_{\chi}(p) \quad \forall p \in [0, 1],$$

satisfying (5).

$$\text{Step 6: } \frac{p_0 - \underline{y}(\tilde{\theta})}{\bar{y}(\tilde{\theta}) - \underline{y}(\tilde{\theta})} \left(\psi_{\chi}(\bar{y}(\tilde{\theta})) + H(\bar{y}(\tilde{\theta}), \tilde{\theta}) \right) + \frac{\bar{y}(\tilde{\theta}) - p_0}{\bar{y}(\tilde{\theta}) - \underline{y}(\tilde{\theta})} \left(\psi_{\chi}(\underline{y}(\tilde{\theta})) + H(\underline{y}(\tilde{\theta}), \tilde{\theta}) \right) - H(p_0, \tilde{\theta}) = 0.$$

We have

$$\begin{aligned}
& \frac{p_0 - \underline{y}(\tilde{\theta})}{\bar{y}(\tilde{\theta}) - \underline{y}(\tilde{\theta})} \left(\psi_{\chi}(\bar{y}(\tilde{\theta})) + H(\bar{y}(\tilde{\theta}), \tilde{\theta}) \right) + \frac{\bar{y}(\tilde{\theta}) - p_0}{\bar{y}(\tilde{\theta}) - \underline{y}(\tilde{\theta})} \left(\psi_{\chi}(\underline{y}(\tilde{\theta})) + H(\underline{y}(\tilde{\theta}), \tilde{\theta}) \right) - H(p_0, \tilde{\theta}) \\
& = \frac{p_0 - \underline{y}(\tilde{\theta})}{\bar{y}(\tilde{\theta}) - \underline{y}(\tilde{\theta})} \left((\bar{y}(\tilde{\theta}) - p_0)H_p(p_0, \tilde{\theta}) \right) + \frac{\bar{y}(\tilde{\theta}) - p_0}{\bar{y}(\tilde{\theta}) - \underline{y}(\tilde{\theta})} \left((\underline{y}(\tilde{\theta}) - p_0)H_p(p_0, \tilde{\theta}) \right) \\
& = 0.
\end{aligned}$$

Thus, ψ_{χ} satisfies (8), and if it is nonnegative, it implements χ by Lemma 8. \square

Proof of Lemma 9: Superregularity Implies Strict Regularity

$\frac{(g'(\theta))^2}{g(\theta)} \left(\frac{1-F(\theta)}{f(\theta)} \right) \geq 0 > g'(\theta)$. Hence, superregularity implies

$$g'(\theta) - \frac{d}{d\theta} \left[\frac{1-F(\theta)}{f(\theta)} g'(\theta) \right] < 0 \Leftrightarrow h''(p) \left(g'(\theta) - \frac{d}{d\theta} \left[\frac{1-F(\theta)}{f(\theta)} g'(\theta) \right] \right) = \frac{\partial}{\partial \theta} \left[H_{pp}(p, \theta) - \frac{1-F(\theta)}{f(\theta)} H_{pp\theta}(p, \theta) \right] > 0.$$

□

Proof of Theorem 7 (Characterization of the Optimal Results-Based Contract) First, I need to show that χ^* is \underline{y}, \bar{y} -smooth. This follows directly if $\Theta_+ = \Theta$. If $\Theta_+ = (\underline{\theta}, 1]$, note that \underline{p}, \bar{p} are continuously differentiable on $(\underline{\theta}, 1]$ and so have Lipschitz extensions \underline{y}, \bar{y} to $[\underline{\theta}, 1]$. These extensions are continuously differentiable: For each $\theta \in (\underline{\theta}, 1]$,

$$\begin{aligned} \pi(\underline{y}(\theta), \theta) + \pi_p(\underline{y}(\theta), \theta)(\bar{y}(\theta) - \underline{y}(\theta)) - \pi(\bar{y}(\theta), \theta) &= 0, \\ \pi(\underline{y}(\theta), \theta) + \pi_p(\bar{y}(\theta), \theta)(\bar{y}(\theta) - \underline{y}(\theta)) - \pi(\bar{y}(\theta), \theta) &= 0 \end{aligned}$$

By continuity, this holds at $\underline{\theta}$. Since $\bar{p}(\underline{\theta}) - \underline{p}(\underline{\theta}) \neq 0$ from Lemma 7, it follows from Lemma 12 that the extensions \underline{y}, \bar{y} are continuously differentiable. Thus, χ^* is \underline{y}, \bar{y} -smooth.

(i): Strict regularity follows from Lemma 9 and superregularity. The result then follows from Theorem 4.

(v): From the proof of Theorem 6, we have

$$\psi''_{\chi^*}(p) = \begin{cases} 0, & p \in [y(1), 1), \\ -H_{pp}(p, \bar{p}^{-1}(p)) + \frac{-\int_{\underline{p}(\bar{p}^{-1}(p))}^p H_{pp\theta}(r, \bar{p}^{-1}(p)) dr dz}{(p - \underline{p}(\bar{p}^{-1}(p))) \bar{p}'(\bar{p}^{-1}(p))}, & p \in (\bar{p}(\underline{\theta}), \bar{p}(1)), \\ 0, & p \in (\underline{p}(\underline{\theta}), \bar{p}(\underline{\theta})), \\ -H_{pp}(p, \underline{p}^{-1}(p)) + \frac{\int_{\bar{p}(\underline{p}^{-1}(p))}^p H_{pp\theta}(r, \underline{p}^{-1}(p)) dr dz}{(\bar{p}(\underline{p}^{-1}(p)) - p) \underline{p}'(\underline{p}^{-1}(p))}, & p \in (\underline{p}(1), \underline{p}(\underline{\theta})), \\ 0, & p \in (0, \underline{p}(1)]. \end{cases}$$

From Lemma 12 and Theorem 4 we have

$$\begin{aligned} \frac{dp}{d\theta} &= \frac{\int_{\underline{p}(\theta)}^{\bar{p}(\theta)} \int_{\underline{p}(\theta)}^z \frac{\partial}{\partial \theta} \left(H_{pp}(r, \theta) - \frac{1-F(\theta)}{f(\theta)} H_{pp\theta}(r, \theta) \right) dr dz}{\left(w''(\underline{p}(\theta)) + H_{pp}(\underline{p}(\theta), \theta) - \frac{1-F(\theta)}{f(\theta)} H_{pp\theta}(\underline{p}(\theta), \theta) \right) (\bar{p}(\theta) - \underline{p}(\theta))}, \\ \frac{d\bar{p}}{d\theta} &= \frac{-\int_{\underline{p}(\theta)}^{\bar{p}(\theta)} \int_z^{\bar{p}(\theta)} \frac{\partial}{\partial \theta} \left(H_{pp}(r, \theta) - \frac{1-F(\theta)}{f(\theta)} H_{pp\theta}(r, \theta) \right) dr dz}{\left(w''(\bar{p}(\theta)) + H_{pp}(\bar{p}(\theta), \theta) - \frac{1-F(\theta)}{f(\theta)} H_{pp\theta}(\bar{p}(\theta), \theta) \right) (\bar{p}(\theta) - \underline{p}(\theta))}, \end{aligned}$$

where \underline{p}^{-1} and \bar{p}^{-1} denote the inverses of \underline{p} and \bar{p} , respectively. Thus $\psi''_{\chi^*}(p)$ is given by

$$\begin{aligned} &0, \quad p \in [y(1), 1), \\ -H_{pp}(p, \bar{p}^{-1}(p)) + \frac{\left(w''(p) + H_{pp}(p, \bar{p}^{-1}(p)) - \frac{1-F(\bar{p}^{-1}(p))}{f(\bar{p}^{-1}(p))} H_{pp\theta}(p, \bar{p}^{-1}(p)) \right) \int_{\underline{p}(\bar{p}^{-1}(p))}^p \int_z^{\bar{p}^{-1}(p)} H_{pp\theta}(r, \bar{p}^{-1}(p)) dr dz}{\int_{\underline{p}(\bar{p}^{-1}(p))}^p \int_z^{\bar{p}^{-1}(p)} \frac{\partial}{\partial \theta} \left[H_{pp}(r, \theta) - \frac{1-F(\theta)}{f(\theta)} H_{pp\theta}(r, \theta) \right]_{\theta=\bar{p}^{-1}(p)} dr dz}, & p \in (\bar{p}(\underline{\theta}), \bar{p}(1)), \\ &0, \quad p \in (\underline{p}(\underline{\theta}), \bar{p}(\underline{\theta})), \\ -H_{pp}(p, \underline{p}^{-1}(p)) + \frac{\left(w''(p) + H_{pp}(p, \underline{p}^{-1}(p)) - \frac{1-F(\underline{p}^{-1}(p))}{f(\underline{p}^{-1}(p))} H_{pp\theta}(p, \underline{p}^{-1}(p)) \right) \int_{\bar{p}(\underline{p}^{-1}(p))}^p \int_z^{\underline{p}^{-1}(p)} H_{pp\theta}(r, \underline{p}^{-1}(p)) dr dz}{\int_{\bar{p}(\underline{p}^{-1}(p))}^p \int_z^{\underline{p}^{-1}(p)} \frac{\partial}{\partial \theta} \left[H_{pp}(r, \theta) - \frac{1-F(\theta)}{f(\theta)} H_{pp\theta}(r, \theta) \right]_{\theta=\underline{p}^{-1}(p)} dr dz}, & p \in (\underline{p}(1), \underline{p}(\underline{\theta})), \\ &0, \quad p \in (0, \underline{p}(1)], \end{aligned}$$

or equivalently

$$\begin{aligned}
& 0, \quad p \in [y(1), 1), \\
& \frac{\int_{\underline{p}(\bar{p}^{-1}(p))}^{\underline{p}} \int_z^{\underline{p}} H_{pp\theta}(r, \bar{p}^{-1}(p)) \left(w''(p) - \frac{1-F(\bar{p}^{-1}(p))}{f(\bar{p}^{-1}(p))} H_{pp\theta}(p, \bar{p}^{-1}(p)) \right) + H_{pp}(p, \bar{p}^{-1}(p)) \left(\frac{\partial}{\partial \theta} \left[\frac{1-F(\theta)}{f(\theta)} H_{pp\theta}(r, \theta) \right]_{\theta=\bar{p}^{-1}(p)} \right) dr dz}{\int_{\underline{p}(\bar{p}^{-1}(p))}^{\underline{p}} \int_z^{\underline{p}} \frac{\partial}{\partial \theta} \left[H_{pp}(r, \theta) - \frac{1-F(\theta)}{f(\theta)} H_{pp\theta}(r, \theta) \right]_{\theta=\bar{p}^{-1}(p)} dr dz}, \quad p \in (\bar{p}(\underline{\theta}), \bar{p}(1)), \\
& 0, \quad p \in (\underline{p}(\underline{\theta}), \bar{p}(\underline{\theta})), \\
& \frac{\int_{\underline{p}}^{\bar{p}(\underline{p}^{-1}(p))} \int_p^z H_{pp\theta}(r, \underline{p}^{-1}(p)) \left(w''(p) - \frac{1-F(\underline{p}^{-1}(p))}{f(\underline{p}^{-1}(p))} H_{pp\theta}(p, \underline{p}^{-1}(p)) \right) + H_{pp}(p, \underline{p}^{-1}(p)) \left(\frac{\partial}{\partial \theta} \left[\frac{1-F(\theta)}{f(\theta)} H_{pp\theta}(r, \theta) \right]_{\theta=\underline{p}^{-1}(p)} \right) dr dz}{\int_{\underline{p}}^{\bar{p}(\underline{p}^{-1}(p))} \int_p^z \frac{\partial}{\partial \theta} \left[H_{pp}(r, \theta) - \frac{1-F(\theta)}{f(\theta)} H_{pp\theta}(r, \theta) \right]_{\theta=\underline{p}^{-1}(p)} dr dz}, \quad p \in (\underline{p}(1), \underline{p}(\underline{\theta})), \\
& 0, \quad p \in (0, \underline{p}(1)].
\end{aligned}$$

Consider the expressions for $\psi''_{\chi^*}(p)$ on the regions where it is nonzero. The integrand in their denominators is positive by strict regularity, so $\psi''_{\chi^*}(p) \geq 0$ everywhere if and only if the numerators are positive. Multiplicative separability allows the integrands in the numerators to be written as (writing $\underline{p}^{-1}(p)$ and $\bar{p}^{-1}(p)$ as θ for simplicity)

$$g'(\theta)h''(r)w''(p) + h''(r)h''(p) \left[\frac{\partial}{\partial \theta} \left[\frac{1-F(\theta)}{f(\theta)} g'(\theta) \right] g(\theta) - (g'(\theta))^2 \frac{1-F(\theta)}{f(\theta)} \right].$$

Since $g'(\theta) < 0 \forall \theta$, $h''(p) < 0 \forall p$, and $w''(p) \geq 0 \forall p$, showing that the part in brackets is positive for all θ suffices to show that the entire expression is. This follows directly from superregularity.

(iv): By (v), it suffices to show that $\psi'_{\chi^*}(\underline{p}(\underline{\theta})) \leq 0$ and $\psi'_{\chi^*}(\bar{p}(\underline{\theta})) \geq 0$. We have

$$\begin{aligned}
\psi'_{\chi^*}(\bar{p}(\underline{\theta})) &= H_p(p_0, \underline{\theta}) - H_p(\bar{p}(\underline{\theta}), \underline{\theta}) \geq 0, \\
\psi'_{\chi^*}(\underline{p}(\underline{\theta})) &= H_p(p_0, \underline{\theta}) - H_p(\underline{p}(\underline{\theta}), \underline{\theta}) \leq 0,
\end{aligned}$$

by concavity of H , as desired.

(iii): By (iv), it suffices to show that $\psi_{\chi^*}(\underline{p}(\underline{\theta})) \geq 0$ and $\psi_{\chi^*}(\bar{p}(\underline{\theta})) \geq 0$. We have

$$\begin{aligned}
\psi_{\chi^*}(\bar{p}(\underline{\theta})) &= H(p_0, \underline{\theta}) - H(\bar{p}(\underline{\theta}), \underline{\theta}) + (\bar{p}(\underline{\theta}) - p_0)H_p(p_0, \underline{\theta}) \\
&= - \int_{p_0}^{\bar{p}(\underline{\theta})} \int_{p_0}^z H_{pp}(r, \underline{\theta}) dr dz \geq 0, \\
\psi_{\chi^*}(\underline{p}(\underline{\theta})) &= H(p_0, \underline{\theta}) - H(\underline{p}(\underline{\theta}), \underline{\theta}) + (\underline{p}(\underline{\theta}) - p_0)H_p(p_0, \underline{\theta}) \\
&= \int_{\underline{p}(\underline{\theta})}^{p_0} \int_z^{p_0} H_{pp}(r, \underline{\theta}) dr dz \geq 0.
\end{aligned}$$

(ii): Follows immediately from Theorem 6. □