

Cambridge Working Paper Economics

Cambridge Working Paper Economics: 1712

Hidden Testing and Selective Disclosure of Evidence

Claudia
Herresthal

11 December 2017

(First submitted to the CWPE Series 19 February 2017)

This paper contrasts a decision maker's payoff under public and private information acquisition by a biased advisor. Both players agree on the optimal choice under certainty, but differ in how they trade off the loss from errors. The advisor can sequentially acquire informative test outcomes. If acquisition is private he decides in the final period which realizations to verifiably disclose. If players' preferences are sufficiently misaligned, the decision maker is weakly better off under private rather than public information acquisition. The effect on the advisor's payoff depends on the direction of his bias.

Hidden Testing and Selective Disclosure of Evidence*

Claudia Herresthal[†]

December 11, 2017

(First submitted to the CWPE Series 19 February 2017)

Abstract

This paper contrasts a decision maker's payoff under public and private information acquisition by a biased advisor. Both players agree on the optimal choice under certainty, but differ in how they trade off the loss from errors. The advisor can sequentially acquire informative test outcomes. If acquisition is private he decides in the final period which realizations to verifiably disclose. If players' preferences are sufficiently misaligned, the decision maker is weakly better off under private rather than public information acquisition. The effect on the advisor's payoff depends on the direction of his bias.

Keywords: endogenous information acquisition, verifiable disclosure, transparency

JEL codes: D83, D82

*I would like to thank Meg Meyer, as well as Inga Deimen, Matt Elliott, Robert Evans, Pawel Gola, Ole Jann, Arina Nikandrova, Anna Sanktjohanser and the conference and seminar participants at Conference on Economic Design, Game Theory Conference at Stony Brook, EEA, Bocconi, Birkbeck, Oxford and Cambridge for valuable comments and feedback.

[†]Queens' College, University of Cambridge, Silver Street, Cambridge CB3 9ET, UK, cush2@cam.ac.uk.

1 Introduction

Pharmaceutical companies have recently come under scrutiny for selectively reporting outcomes of clinical trials. As a response to these pressures, several of these companies have pledged to register all of their trials in open online databases and to publish their findings (Goldacre et al. (2017)). At first sight, it seems that demanding such transparency would improve regulation, because pharmaceutical companies can no longer omit trials with unfavorable outcomes. However, companies may also strategically respond by changing which trials they run and this could leave the regulator to base his approval decision on weaker evidence.

This paper studies when and why a decision maker (e.g. a regulator) prefers hidden to observable information acquisition by a biased advisor (e.g. a company). The decision maker has to either accept or reject a given hypothesis. Both agree on the optimal action under certainty, but when there is some uncertainty about whether the hypothesis is true or false they may trade off the mistake from inappropriate actions differently. For example, both the regulator and the company may agree that a drug should be approved only if it is safe, but the regulator is relatively more averse to the mistake of approving an unsafe drug than the company. I first consider a setting of *hidden testing*, which corresponds to not having a trial registry. The advisor can in private sequentially acquire binary test outcomes over a fixed number of periods.¹ In the final period, the advisor chooses which outcomes to disclose and then the decision maker takes an action. The advisor can verifiably disclose outcomes, i.e. he can omit but not fake outcomes, but he cannot verifiably disclose the period in which these outcomes were discovered. In the example, this represents the fact that the regulator does not know how many trials have been run in total by the time he receives an application for approval. I contrast this with a setting of *observable testing*, which corresponds to having a trial registry. When testing is observable, the advisor has no private information.²

My main focus lies on situations in which the decision maker (DM) cannot commit to a decision rule, e.g. if on the basis of evidence already presented the regulator prefers approval to rejection, then he has limited scope to delay approval unless additional positive evidence can be produced.³ I first study a two-period model and then show that the main insights are robust to allowing for an arbitrary fixed number of periods. My main result fully characterizes the situations in which the DM is strictly better off and the situations in which she is strictly worse off under hidden rather than observable testing. I identify two distinct economic effects which cause the DM to benefit from hidden testing, and I refer to them as the skepticism effect and the insurance effect.

¹The finite horizon represents the fact that companies may face time and money constraints in their development. This can be seen as a limiting case of a convex cost of testing, as discussed in Section 5.

²For my results, the important aspect of observable testing is that all realizations are disclosed. It is irrelevant whether or not the time at which these outcomes were discovered is disclosed.

³Although a regulator can demand additional evidence, he may find it costly to delay approval of a drug, e.g. because he is facing pressure by patient advocacy groups, as for example in the case of the drug Flibanserin (Pollack (2015)) or drugs for HIV treatment (Epstein (1998)).

The skepticism effect describes situations in which the DM can credibly set a tougher standard for acceptance in equilibrium under hidden than under observable testing, which allows her to take a more informed decision. In a two-period model, suppose the DM wants to accept if and only if there are strictly more positive than negative outcomes and the advisor is more inclined to accept than the DM. Under observable testing, the advisor stops testing as soon as he has found one positive outcome and the DM accepts. However, if the advisor were to reveal one positive outcome under hidden testing, the DM may reject. The reason is that the DM suspects that the advisor may have run a second test and omitted contradicting evidence. If the DM is sufficiently averse to falsely accepting, the advisor needs to show two positive outcomes for her to accept. In the example, if trials are registered then fewer trials with positive findings may be needed to convince the regulator. Although the evidence taken at face-value is stronger when trials are registered, the regulator also lowers his threshold for how much positive evidence is needed for approval. The skepticism effect shows that this can lead to less safe drugs appearing on the market.

The insurance effect describes situations in which the DM gains better information when testing is hidden, because the advisor has greater incentives to learn whether or not the hypothesis is true. In a two-period model, suppose the DM wants to reject if and only if at least one in two tests is negative, whereas the advisor wants to reject if and only if both are negative.⁴ Under observable testing, the advisor runs no test at all and the DM accepts. To see why, suppose the advisor were to find a negative outcome. Then the DM would reject and she would reject irrespective of what a second test yields. However, a single negative outcome is not sufficient for the advisor to prefer rejection. By not testing, the advisor can guarantee that the DM acts in his interest. By contrast, under hidden testing, the advisor can test in private and hide outcomes if he wants to accept, but reveal all outcomes if he wants to reject. Therefore, the DM rejects if and only if the evidence is sufficiently negative to convince the advisor to reject. The insurance effect shows that a trial registry may deter companies from investigating whether or not certain side effects exist. To see why, suppose the regulator would approve the drug in the absence of further information. Without a registry, the company may run investigations in private and withdraw its application if it finds strong evidence that these side effects exist. However, with a registry, it may not conduct any investigations out of fear that it will find some evidence of these side effects, which is too weak for the company to lose interest in selling the drug, but strong enough for the regulator to decline approval.

However, the DM does not necessarily benefit from hidden testing. Although the advisor becomes better informed when testing is hidden, he also has a larger scope for strategic disclosure because he has acquired more test outcomes. I show that if players' preferences are sufficiently

⁴In this example, I assume that the test technology is such that a false positive outcome is more likely to arise than a false negative outcome. In a model with more than two periods, the insurance effect can exist even when this assumption does not hold.

misaligned then the DM weakly benefits from hidden testing. This demonstrates that whether or not a trial registry is optimal may depend on how much the regulator values the drug relative to the pharmaceutical company. For example, the regulator may be strongly averse to approving an unsafe drug relative to the company when a drug is intended mostly for cosmetic use, and then it would be better if the company selectively disclosed trials. Overall, hidden testing allows the DM to achieve her first-best expected payoff for a larger set of parameter combinations than observable testing, where her first-best expected payoff is defined as the one achieved if the DM could acquire evidence herself.

If the DM had the power to commit to her decision rule, she can achieve her first-best expected payoff, whether testing is hidden or observable. Therefore, the DM would ideally like to commit to a decision rule, but if this option is not available to her, then she may compensate for this to some extent by allowing the advisor to test in private. In addition, I show that if the advisor has the power to commit to a testing and a disclosure plan, then the skepticism effect ceases to exist, whereas the insurance effect persists. This provides a rationale for why pharmaceutical companies faced demands to register their trials for any drug, thereby precluding that companies could choose to side-step registries for individual drugs. In addition, I show that delegating decision rights to the advisor never makes the DM strictly better off than hidden testing. This implies that it is not necessary to give up decision rights in order to improve the advisor's incentives for information acquisition provided this information acquisition is private.

Although I use the drug approval process as my leading example, the insights also apply to the legislative process. By law, both the prosecutor and the judge should seek conviction if and only if the defendant is guilty. However, if the evidence is inconclusive, the prosecutor is more eager to convict than the judge. The US supreme court ruled that a prosecutor suppressing exculpatory evidence is violating due process (*Brady v. Maryland*, 1963). Yet, charges for prosecutorial misconduct are rare (Davis (2006)). It may be optimal not to deter the prosecutor from hiding evidence. This is because the prosecutor then has to provide stronger evidence that the defendant is guilty compared to a situation in which prosecutors are deterred from hiding evidence (skepticism effect). In addition, if the prosecutor is deterred from hiding evidence, then he may not pursue certain lines of investigation out of fear that the evidence in favor of the defendant's innocence could be weak, yet lead the judge to acquit (lack of insurance effect). Despite the low rate of charges for prosecutorial misconduct in general, they constitute the reason for a large share of reversals of capital convictions (Liebman et al. (2000)). For capital cases, both the judge and the prosecutor may face a more similar trade-off since the stakes of convicting if innocent are high and, therefore, deterring prosecutors from hiding evidence is desirable.

Another example is scientific research. Scientists care about informing the public, but they are also under pressure to publish and therefore may be less averse to accepting a false hypothesis than the public. Recently, the fact that the results of many scientific studies cannot be replicated has

strengthened demands for pre-analysis plans.⁵ However, if scientists have to register their experiment then it is harder for an editor to be credibly skeptical of the significance of their findings and, hence, harder for him to demand additional robustness checks before accepting a paper (skepticism effect). In addition, if scientists have to register their experiment then they may not include certain aspects in their research agenda out of fear that the additional findings might cast doubt on their existing conclusions and, hence, lead the editor to reject their work, whereas the scientists would still want to see their conclusions published (lack of insurance effect).

The existing literature has compared a DM's payoff under hidden and observable information acquisition, e.g. Brocas and Carrillo (2007), Henry (2009) and Felgenhauer and Loerke (2017). Compared to these papers, my framework allows for the conflict of interest between players to be endogenously determined. Put differently, whether players agree or disagree on the optimal action depends on what evidence has been acquired. In contrast, these papers assume either that the advisor's optimal action differs from the DM's optimal action by a constant in any given state (Henry (2009)) or that the advisor has state-independent preferences (Brocas and Carrillo (2007), Felgenhauer and Loerke (2017)).⁶ Allowing for an endogenous conflict of interest enables me to identify the insurance effect as an explanation of why hidden testing can benefit the DM. This effect seems intuitive, yet has not been highlighted in the existing literature.

The key distinction in predictions between my model and the related literature is that the DM can either be strictly worse or strictly better off under hidden testing, depending on how closely aligned players' preferences are. This matches our intuition that in some applications hidden testing can be harmful and a greater push towards transparency could improve decision making. By contrast, the related literature predicts that hidden testing is weakly better than observable testing. The contrast with their findings is driven by the combination of assumptions I make. In particular, I assume that the advisor cannot verifiably disclose his beliefs, the DM's choice is discrete and the testing technology is fixed. In many real-life situations, an advisor would have some flexibility with regard to the test he runs, but not full flexibility. Felgenhauer and Loerke (2017) study one extreme, i.e. the advisor has full flexibility with regard to designing the test, while I focus on the other extreme, i.e. the test technology is given. Comparing our findings illustrates that whether the DM always benefits from hidden testing is sensitive to how much discretion the advisor has with regard to the tests he designs. A more detailed discussion of these closely related papers and the extensive literature can be found in Section 6 after the analysis has been completed to allow for an informed comparison.

The paper is structured as follows. Section 2 introduces the model. Section 3 contains the key comparison of the DM's expected payoff under hidden versus observable testing first for the two-

⁵E.g. see the Berkeley Initiative for Transparency in the Social Sciences.

⁶Such state-independent preferences are a special case of my framework when the advisor's loss from falsely accepting is 0.

period case and then for the case with any arbitrary finite number of periods. Section 4 analyzes this comparison when the DM delegates decision rights to the advisor and when one of the players has commitment power. Finally, Section 5 analyzes the comparison when the advisor has to commit to the number of tests ex ante and when the horizon is infinite and the advisor incurs a constant cost per test. Section 6 presents a detailed comparison with the related literature. Section 7 concludes. All proofs can be found in the Appendix.

2 Model

I will first outline the model of hidden testing and then outline the benchmark model of observable testing.

Setting: Time is discrete and there are finitely many periods, $n = 1, \dots, N$. The players are a Decision Maker (DM) and an advisor (A).

Timing and Information: In period $n = 1$, Nature draws a state $\omega \in \{false, true\}$, which determines whether a given hypothesis is true or false, where $Pr(true) = q \in [0, 1]$. The realized state is unobserved by both players. In each period $n = 1, \dots, N$, the advisor first privately chooses whether to test $a = 1$ or not $a = 0$. If he tests then Nature draws a test outcome $s_n \in \{+, -\}$, with state-dependent accuracy given by

$$Pr(-|false) = p_F \tag{1}$$

$$Pr(+|true) = p_T, \tag{2}$$

where $\frac{1}{2} < p_i < 1$ for $i = F, T$ and the realizations of outcomes are independent conditional on the state.⁷ If the advisor does not test then Nature does not draw a test outcome, and this event is denoted by $s_n = \emptyset$. The test outcome is privately observed by the advisor. At the end of period N , the advisor sends a message $m \in M$ to the DM, where the message space M is defined below. Then the DM chooses an action $\tau \in \{accept, reject\}$. Finally, payoffs are realized. It is assumed that players do not have commitment power.

Payoffs: A player $i = DM, A$ incurs the following state-dependent loss, denoted by $L_i(\tau, \omega)$:⁸

$L_i(\tau, \omega)$	$false$	$true$
$reject$	0	1
$accept$	λ_i	0

where $\lambda_i \geq 0$. Hence, if the state of the world were known then players would agree on what the optimal action is. Each player would prefer to accept if and only if the hypothesis is true. However,

⁷In Section 3.1, I allow for the informativeness of the test to depend on the state of the world, i.e. $p_T \neq p_F$.

⁸This framework is based on DeGroot (1970)'s framework for optimal statistical decisions.

when the state of the world is uncertain then players may disagree on what the optimal action is, because players may face a different trade-off between the loss from accepting if false and the loss from rejecting if true. Without loss of generality, I normalize the loss from rejecting if true to be 1 for each player, but allow for the loss from accepting if false to differ between players. The advisor is *more reluctant* to accept if he cares more about the loss from falsely accepting than the DM, i.e. if $\lambda_A \geq \lambda_{DM}$. Conversely, the advisor is *more enthusiastic* about accepting if he cares less about the loss from falsely rejecting than the DM, i.e. if $\lambda_A \leq \lambda_{DM}$.

Player i 's expected payoff is given by $\pi_i(\tau) = -E(L_i(\tau))$, and the vector of expected payoffs is given by $\Pi = (\pi_A, \pi_{DM})$. Each player i maximizes their expected payoff π_i arising from the statistical decision problem.⁹ Without loss of generality, I assume the DM's loss λ_{DM} from falsely accepting is sufficiently high such that her expected payoff at the prior belief q is maximized when choosing *reject*, i.e. *reject* is her default choice:¹⁰

$$\pi_{DM}(\text{reject}) > \pi_{DM}(\text{accept}) \Leftrightarrow \frac{q}{1-q} < \lambda_{DM}.$$

Histories: Let the *history of test outcomes* at the end of period $n = 1, \dots, N$ be denoted by $h_n \in H_n$, which is an ordered list of past draws by Nature, i.e. $h_n = (s_1, s_2, \dots, s_n)$ where $s_n \in \{+, -, \emptyset\}$. Let h_0 denote the history of outcomes at the start of the game. Let the *unordered history of test outcomes* at the end of period N be denoted by \tilde{h} , and the set of all \tilde{h} be denoted by \tilde{H} . The unordered history \tilde{h} is a set of past draws by Nature, which contains the realization of past draws, but not the order in which they were obtained. For example, if $N = 2$ then both $h_2 = (+, -)$ and $h_2 = (-, +)$ map into $\tilde{h} = \{+, -\}$.

Message: The advisor's message space is given by $M = \mathcal{P}(\tilde{h}) \in \mathcal{P}(\tilde{H})$, where $\mathcal{P}(\tilde{h})$ denotes the power set of \tilde{h} , i.e. the set of all subsets of \tilde{h} including the empty set. For example, if $N = 2$ and $\tilde{h} = \{+, -\}$ then the message space is given by $M = \{\{\emptyset\}, \{+\}, \{-\}, \{+, -\}\}$. The advisor can report any subset of test outcomes he has generated, which implies that he cannot make up outcomes but he is able to hide outcomes (verifiable disclosure). In addition, it implies that the advisor cannot verifiably disclose the period in which the outcome was found (not datable), i.e. if $N = 2$ and he reports $m = \{+\}$ he cannot convey whether $h_2 = (+, \cdot)$ or $h_2 = (\cdot, +)$.

Strategies - The advisor has a testing and a disclosure strategy. A testing strategy for the advisor is: $\sigma_A : H_n \rightarrow \{0, 1\}$. It selects action $a \in \{0, 1\}$ conditional on history $h_n \in H_n$ for $n = 0, \dots, N - 1$. A reporting strategy for the advisor is $\sigma_M : H_N \rightarrow M$. It selects a message $m \in M$ conditional on history $h_N \in H_N$. A strategy for the DM is $\sigma_{DM} : M \rightarrow T$ where $T \in \{\text{accept}, \text{reject}\}$. It selects action $\tau \in T$ conditional on message $m \in M$.

⁹I assume tests are costless for tractability, i.e. the advisor stops testing if and only if he can strategically manipulate the DM's choice. All my results continue to hold if the cost of testing is sufficiently small.

¹⁰If I restricted preferences such that *accept* is the DM's default choice, the results would be a "mirror image" of the results presented here.

Equilibrium Concept: The solution concept is a sequential equilibrium in pure strategies.¹¹ A sequential equilibrium consists of both a profile of strategies $\sigma \equiv (\sigma_A, \sigma_M, \sigma_{DM})$ and a system of beliefs $\mu : \cup_{n=1}^N H_n \cup M \rightarrow \Delta\Omega$, where $\Omega = \{true, false\}$. Beliefs select a probability distribution over states $\omega \in \Omega$ for each history of outcomes $h_n \in H_n$ where $n = 1, \dots, N$ and for each message $m \in M$. By definition of a sequential equilibrium,

1. the advisor's strategy (σ_A, σ_M) at any history h_n for $n = 0, \dots, N - 1$ maximizes his expected payoff given the DM's strategy σ_{DM} and given the system of beliefs μ , and
2. the DM's strategy σ_{DM} maximizes her expected payoff at any message m given the system of beliefs μ , and
3. there exists a sequence of completely mixed strategies $\{\sigma^k\}_{k=1}^\infty$, with $\lim_{k \rightarrow \infty} \sigma^k = \sigma$, such that the system of beliefs $\mu = \lim_{k \rightarrow \infty} \mu^k$, where μ^k denotes the beliefs derived from strategy profile σ^k using Bayes' rule.

In case of multiple such equilibria, I select an equilibrium in which the advisor achieves the highest expected payoff, i.e. an *advisor-preferred equilibrium*.

Observable Testing: In the benchmark case of observable testing, the advisor's actions and Nature's draws of test outcomes are observed by both players. This renders the advisor's report superfluous. The DM's strategy is given by $\bar{\sigma}_{DM} : H_N \rightarrow T$. It selects action $\tau \in T$ conditional on history $h_N \in H_N$ (rather than conditional on a message). A sequential equilibrium consists of a strategy profile $\bar{\sigma} \equiv (\bar{\sigma}_A, \bar{\sigma}_{DM})$ and a system of beliefs $\bar{\mu} : \cup_{n=1}^N H_n \rightarrow \Delta\Omega$, where $\Omega = \{true, false\}$. By definition of a sequential equilibrium:

1. the advisor's strategy $\bar{\sigma}_A$ at any history h_n for $n = 0, \dots, N$ maximizes his expected payoff given the DM's strategy $\bar{\sigma}_{DM}$ and given the system of beliefs $\bar{\mu}$, and
2. the DM's strategy $\bar{\sigma}_{DM}$ maximizes her expected payoff at any any history h_N given the system of beliefs $\bar{\mu}$, and
3. there exists a sequence of completely mixed strategies $\{\bar{\sigma}^k\}_{k=1}^\infty$, with $\lim_{k \rightarrow \infty} \bar{\sigma}^k = \bar{\sigma}$, such that the system of beliefs $\bar{\mu} = \lim_{k \rightarrow \infty} \bar{\mu}^k$, where $\bar{\mu}^k$ denotes the beliefs derived from strategy profile $\bar{\sigma}^k$ using Bayes' rule.

Remarks on Equilibrium Selection: I choose an advisor-preferred equilibrium to uniquely identify expected equilibrium payoffs. This selection allows me to demonstrate that even if the advisor fully exploits the additional discretion he is given under hidden testing, the DM can still be strictly better off under hidden testing. In addition, if the DM is better off under hidden testing

¹¹The restriction to pure strategies is for tractability. The main insights of Theorem 1 (for $N = 2$) and Proposition 3 (for $N > 2$) would apply if mixed strategies are allowed, but the boundaries of regions would change.

in this equilibrium, then she is better off under hidden testing in any equilibrium that is not Pareto dominated.

Remarks on Modeling Choices: I assume the content of evidence is verifiable because I want to focus on the effect of strategically omitting evidence. As the literature on verifiable disclosure has frequently argued, in many settings it is relatively harder to fabricate than to omit evidence and this model assumes an extreme case in which fabrication is prohibitively costly. In addition, I assume that the time at which evidence was collected cannot be verifiably disclosed and the advisor cannot report before the final period. These assumptions are made in order to avoid that the DM can make any inference about the number of tests conducted overall. For example, in a research context, it is often unknown how many experiments a scientist has carried out before he presents his findings, because it is unknown at what point the scientist started his inquiry, or because the time an experiment lasts is unknown or random, or because it is unclear whether or not the scientist had to devote time to other projects. The model has a finite horizon in order to incorporate the fact that the advisor's choice to continue testing depends not only on his beliefs about the state, but also on how many tests he has already run. A finite horizon can be seen as a tractable limiting case of a convex cost of testing, as discussed in Section 5. The DM herself cannot run any tests because I assume that she lacks the time or expertise to do so, e.g. she is not trained in scientific methodologies.

3 Hidden vs. Observable Testing

This section compares the DM's expected payoff under hidden and observable testing. For the case $N = 2$, I fully characterize under which conditions the DM is strictly better or strictly worse off under hidden relative to observable testing. I identify two effects for why the DM's expected payoff improves with hidden testing. Then I give sufficient conditions for these effects to exist when $N > 2$. For any N , I show that the DM is weakly better off under observable testing when preferences are sufficiently misaligned.

3.1 Two-period Model

Throughout this subsection, I assume that $N = 2$. Each player's expected payoff is unique given my equilibrium selection.

Lemma 1 (Equilibrium) *Equilibrium exists under both observable and hidden testing.*

1. *Under observable testing, any equilibrium yields the same expected payoff vector.*
2. *Under hidden testing, any advisor-preferred equilibrium yields the same expected payoff vector.*

To compare expected payoffs in equilibrium, it is useful to consider a modification of game in which Nature draws both test outcomes ex ante and the advisor learns about an outcome if and only if he chooses to test in the corresponding period. Under observable testing, any equilibrium leads to the same mapping from Nature's draws of outcomes to the DM's actions. Therefore, the equilibrium payoff vector is unique. Under hidden testing, the mapping from Nature's draws to the DM's actions is not necessarily unique in equilibrium. However, the mapping that maximizes the advisor's expected payoff is unique. Selecting an advisor-preferred equilibrium when testing is hidden has the following implication. If the DM achieves a higher expected payoff in an advisor-preferred equilibrium under hidden testing than in equilibrium under observable testing, she is better off under hidden rather than observable testing in any equilibrium that is not Pareto dominated.

Characterizations of the equilibrium under both observable and hidden testing can be found in the Appendix in Sections A.1 and A.2. Here I will only illustrate what each player's preferred action $\tau \in \{accept, reject\}$ is given they observe a history of test outcomes h_N . In equilibrium, a player's posterior beliefs about whether the hypothesis is true or not will depend only on the total number of positive and negative outcomes found, i.e. it will depend only on the unordered history of outcomes \tilde{h} . A player i optimally chooses to accept if and only if

$$\begin{aligned} \pi_i \left(accept | \tilde{h} \right) = -Pr \left(false | \tilde{h} \right) \lambda_i \geq -Pr \left(true | \tilde{h} \right) &= \pi_i \left(reject | \tilde{h} \right) \\ \Leftrightarrow \\ \lambda_i \leq \frac{Pr \left(true | \tilde{h} \right)}{Pr \left(false | \tilde{h} \right)} \equiv x_{\tilde{h}}^*, & \quad (3) \end{aligned}$$

where $x_{\tilde{h}}^*$ denotes the posterior likelihood ratio that the hypothesis is true conditional on the unordered history of outcomes \tilde{h} . Hence, the player chooses to accept if and only if their preference parameter lies above the posterior likelihood ratio $x_{\tilde{h}}^*$.

This illustrates the nature of the conflict of interest between the advisor and the DM, which arises due to different values of the loss from accepting if the hypothesis is false. Suppose the conflict of interest is such that the advisor is more enthusiastic than the DM, i.e. $\lambda_A < \lambda_{DM}$. If the evidence is such that the DM optimally accepts, then the advisor also optimally accepts. This is because if the DM optimally accepts, it must be that $\lambda_{DM} \geq x_{\tilde{h}}^*$, and since $\lambda_A > \lambda_{DM}$, this implies

that $\lambda_A > x_{\bar{h}}$, i.e. the advisor also optimally accepts. In addition, suppose that the advisor is more reluctant than the DM, i.e. $\lambda_A > \lambda_{DM}$. By analogous reasoning, if the evidence is such that the DM optimally rejects, then the advisor also optimally rejects.

To compare the DM's expected payoff between hidden and observable testing, I will define four preference parameter regions, each of which is bounded by posterior likelihood ratios denoted by $x_{\bar{h}}$, as defined in equation (3). Therefore, the boundaries of the regions depend on the prior belief q and the test accuracy (p_F, p_T) . First, let me define what I will refer to as the *skepticism region* to be

$$B_I = \{(\lambda_{DM}, \lambda_A) : \lambda_A \leq x_{\{+,-\}}, \bar{x} < \lambda_{DM} \leq x_{\{+\}}\},$$

where

$$\bar{x} = \frac{q(1 - (1 - p_T)^2)}{(1 - q)(1 - p_F^2)}, \quad (4)$$

and what I will refer to as the *insurance region* to be

$$B_{II} = \{\lambda_{DM}, \lambda_A : x_{\{\emptyset\}} < \lambda_{DM} \leq x_{\{+,-\}}, x_{\{+\}} < \lambda_A \leq x_{\{+,+\}}\}.$$

In addition, let me define two further regions W_I and W_{II} as

$$W_I = \{(\lambda_{DM}, \lambda_A) : \lambda_A \leq x_{\{+,-\}} < \lambda_{DM} < \bar{x}\},$$

and

$$W_{II} = \{\lambda_{DM}, \lambda_A : x_{\{\emptyset\}} < \lambda_{DM} \leq x_{\{+,-\}} < \lambda_A < x_{\{+\}}\}.$$

Theorem 1 (DM Payoff Comparison)

1. The DM is strictly better off under hidden rather than observable testing if and only if
 - (a) preferences lie in the skepticism region, i.e. $(\lambda_{DM}, \lambda_A) \in B_I$, or
 - (b) preferences lie in the insurance region, i.e. $(\lambda_{DM}, \lambda_A) \in B_{II}$.
2. The DM is strictly worse off under hidden rather than observable testing if and only if
 - (a) $(\lambda_{DM}, \lambda_A) \in W_I$, or
 - (b) $(\lambda_{DM}, \lambda_A) \in W_{II}$.

The comparison of the DM's expected payoff is illustrated in Figure 1. There exist parameter combinations such that the DM is strictly better off under hidden rather than observable testing,

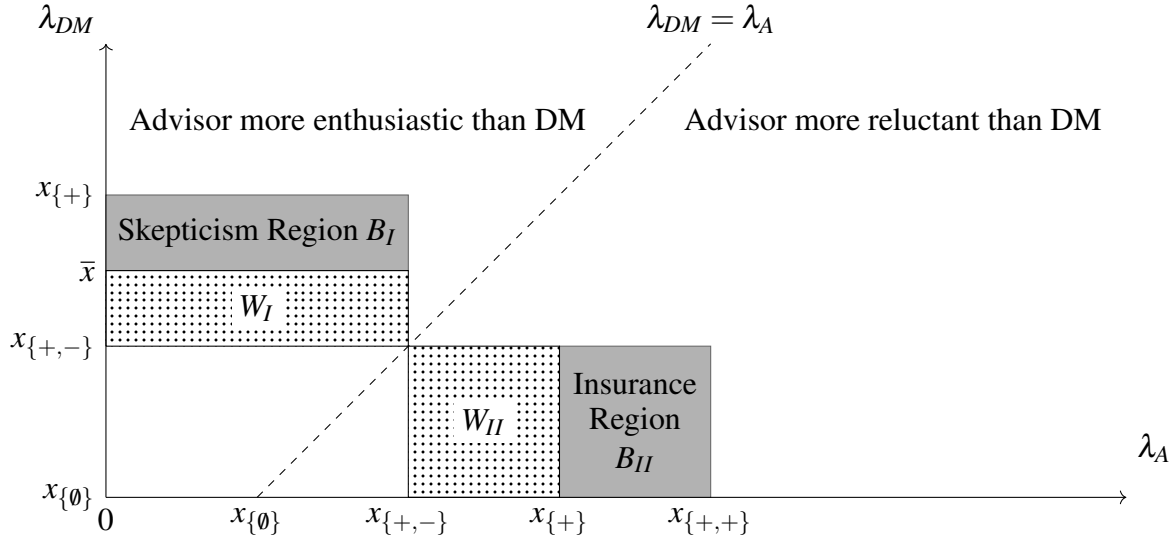


Figure 1: This figure shows the regions in which the DM is strictly better off under hidden rather than observable testing, i.e. the skepticism region B_I and the insurance region B_{II} , and the regions in which the DM is strictly worse off (labeled W_I and W_{II}) in $(\lambda_A, \lambda_{DM})$ -space given $p_T > p_F$. If $p_T \leq p_F$ then regions B_{II} and W_{II} would be empty because there exists no value of λ_{DM} such that $x_{\{\emptyset\}} < \lambda_{DM} \leq x_{\{+, -\}}$.

despite the fact that the advisor enjoys greater discretion when testing is hidden. Both the skepticism region and the insurance region are illustrated by the gray areas in Figure 1. In the skepticism region, the advisor is more enthusiastic about accepting than the DM, i.e. $\lambda_A < \lambda_{DM}$, whereas in the insurance region, the advisor is more reluctant to accept than the DM, i.e. $\lambda_A > \lambda_{DM}$. In what follows, I will explain why the DM benefits from hidden testing in these two regions.

Skepticism Effect Consider preferences in the skepticism region B_I . If the DM knows that a single test was run and the outcome was positive, she accepts. By contrast, if she knows that two tests were run and outcomes were mixed, she rejects. The advisor is more enthusiastic about accepting than the DM, which implies that whenever the DM accepts then both agree that accepting is optimal. Unlike the DM, the advisor would want to accept if he knows two tests were run and the outcomes were mixed.¹²

On the equilibrium path under observable testing, the advisor tests only in the first period and the DM accepts if and only if the outcome is positive. Following a first positive outcome the advisor wants acceptance irrespective of what the second test outcome is. In the absence of any further information the DM acts in the advisor's interest and accepts. Therefore, it is optimal for the advisor to stop the flow of information. The advisor is indifferent between testing or not following a negative first outcome because the DM rejects regardless of what the second outcome

¹²The advisor may even want to accept irrespective of what the test outcomes are.

is.

By contrast, on the equilibrium path under hidden testing, the advisor runs two tests, reveals only positive outcomes to the DM and the DM accepts if and only if both outcomes are positive. It cannot be part of an equilibrium that the advisor reveals a single positive outcome and the DM accepts, as was the case under observable testing. To see why, suppose the DM were to accept based on the report of a single positive outcome. Then the advisor would be tempted to keep testing until he has found a single positive outcome and hide any negative outcomes discovered previously. Therefore, the DM cannot be sure whether the reported positive outcome was found on the first or on the second test. Conditional on the report of a single positive outcome her posterior likelihood ratio is equal to \bar{x} . On balance, the DM finds this too weak evidence in favor of the hypothesis and rejects. However, by revealing two positive outcomes, the advisor can convince the DM to accept.

In conclusion, under hidden testing the DM learns whether a positive outcome is backed up by an additional test, and this allows her to reduce her expected loss. Because the advisor has the temptation to hide negative evidence, the DM can be credibly skeptical towards any report of weak evidence in favor of the hypothesis. As a result, she requires a larger number of positive test outcomes to be convinced.

Insurance Effect Consider preferences in the insurance region B_{II} . This region is non-empty if and only if $p_T > p_F$, i.e. a positive and a negative outcome together are evidence in favor of the hypothesis. The DM rejects without any further information, but if she knows that two tests were run and outcomes were mixed, she accepts. The advisor is more reluctant to accept than the DM, which implies that whenever the advisor accepts then both agree that accepting is optimal. The advisor would want to accept if and only if both outcomes are positive.

On the equilibrium path under observable testing, the advisor does not test at all and the DM rejects. To see why, suppose the advisor were to test in the first period. If the first test outcome was negative, then the DM would reject, but she would reject also in the absence of any test. If the outcome was positive, then the DM would accept and a second test outcome could not change her optimal action. The advisor would approve of the DM's decision if and only if a second test outcome backed the first, but without knowing the outcome of the second test the advisor would want to reject. Therefore, by testing the advisor only has the disadvantage that the evidence could turn out to convince the DM to accept when evidence in favor of the hypothesis is weak.

By contrast, on the equilibrium path under hidden testing, the advisor tests in both periods and reveals all outcomes if and only if both are positive. Otherwise, he reveals nothing. The DM rejects if the advisor reveals nothing and accepts otherwise. The advisor tests because he no longer encounters the risk that the DM does not act in his interest, as he did under observable testing. Instead, he can explore whether or not the evidence is sufficient for him to accept. In case he

is himself convinced that accepting is optimal, he can reveal all his findings, which causes the DM to accept. In case he is not himself convinced he can hide all outcomes. When the advisor reveals nothing, the DM conjectures that the advisor would reject, but she does not know whether or not she would also want to reject. However, her expected posterior belief must equal her prior belief and since her posterior belief that the hypothesis is true is raised when the advisor reveals outcomes, it must be lowered when the advisor reveals nothing. Therefore, the DM optimally rejects when the advisor reveals nothing.

In conclusion, under observable testing, the advisor strategically avoids testing so that he does not risk that the DM accepts when he would want to reject. By allowing the advisor to test in private, the DM insures the advisor against this risk. As a consequence, the DM at least learns whether or not the evidence is strong enough for both to agree to accept.

Although hidden testing can be beneficial for the DM, there are preference parameter regions for which the DM is strictly worse off when testing is hidden. These regions, W_I and W_{II} , are depicted by the dotted areas in Figure 1. It is straightforward to see that preferences are more aligned than in the insurance or skepticism effect regions.

In particular, consider $(\lambda_{DM}, \lambda_A) \in W_I$. The DM has a lower loss from falsely accepting than in the skepticism region B_I . If $(\lambda_{DM}, \lambda_A) \in W_I$, then under observable testing, the advisor tests once and the DM accepts if and only if the outcome is positive, as is the case when $(\lambda_{DM}, \lambda_A) \in B_I$. Under hidden testing, the advisor is again tempted to report a single positive outcome either if a positive outcome was discovered on the first or if it was discovered on the second test. However, in contrast to when $(\lambda_{DM}, \lambda_A) \in B_I$, the DM is willing to accept based on the report of a single positive outcome, since her loss from accepting if false is sufficiently low, i.e. $\lambda_{DM} < \bar{x}$.

In addition, consider $(\lambda_{DM}, \lambda_A) \in W_{II}$. The advisor has a lower loss from falsely accepting compared to in the insurance region B_{II} and, as a consequence, he would accept if he knows that a single test was run and the outcome was positive. Under observable testing, the advisor tests once and the DM accepts if and only if the outcome is positive. The advisor has incentives to test, because after a single positive outcome the DM acts in her interest, unlike in the case when $(\lambda_{DM}, \lambda_A) \in B_{II}$. Under hidden testing, the advisor runs two tests and induces the DM to accept if and only if both outcomes are positive, as is the case when $(\lambda_{DM}, \lambda_A) \in B_{II}$. There is no upside to hidden testing for the DM, because the advisor is already willing to test even if the DM does not insure him against acceptance based on weak evidence.

For any other combination of preferences, the DM is equally well off under hidden as under observable testing, as illustrated by the white areas in Figure 1. If $\lambda_A > x_{\{+,+\}}$ this is because there is no draw of test outcomes such that the advisor would want to accept. Therefore the advisor does not test and the DM rejects, irrespective of whether testing is hidden or observable. If $\lambda_A \leq x_{\{+,+\}}$ then for some regions, there is limited scope for manipulation from a more enthusiastic advisor

due to the fact that the DM does not accept after the first test unless she accepts irrespective of the second test (for $x_{\{+\}} < \lambda_{DM} \leq x_{\{+,+\}}$ and for $\lambda_A < x_{\{\emptyset\}} < \lambda_{DM} \leq x_{\{+,-\}}$). In all other regions, the DM's and the advisor's preferences are sufficiently aligned for the same expected payoff to be achieved under hidden and under observable testing.

Based on the findings of Theorem 1, the DM is weakly better off under hidden rather than observable testing if her preferences are sufficiently misaligned with the advisor's preferences.

Corollary 1 (Preference Alignment) *There exists a threshold $d > 0$ such that the DM is weakly better off under hidden rather than observable testing if*

$$|\lambda_A - \lambda_{DM}| > d.$$

It is also interesting to compare the DM's expected payoff under hidden testing and observable testing to the expected payoff the DM could achieve if she collected the evidence herself. I will refer to the expected payoff in this benchmark case as the DM's *first-best expected payoff*. Let $Z_{OT} \in \mathbb{R}^2$ denote the set of all preference parameters $(\lambda_{DM}, \lambda_A)$ for which the DM achieves her first-best expected payoff under observable testing, and let $Z_{HT} \in \mathbb{R}^2$ denote the set of all preference parameters $(\lambda_{DM}, \lambda_A)$ for which the DM achieves her first-best expected payoff under hidden testing.

Proposition 1 (First-Best Benchmark) *The DM achieves her first-best expected payoff for a larger set of parameter combinations under hidden rather than observable testing, i.e.*

$$Z_{OT} \subset Z_{HT}.$$

The DM's expected payoff differs between hidden and observable testing if and only if preferences lie in the regions B_I , B_{II} , W_I or W_{II} . The DM never achieves her first-best expected payoff under observable testing in any of these regions because the advisor strategically avoids conducting some tests which would be pivotal to the DM's action choice. However, the DM achieves her first-best expected payoff under hidden testing for preferences in the skepticism region B_I . Although she does better under hidden testing if preferences are in the insurance region B_{II} , she does not achieve her first-best expected payoff.

So far I have focused on the changes in the DM's expected payoff. The following proposition characterizes changes in the advisor's expected payoff.

Proposition 2 (Advisor Payoff Comparison)

1. *The advisor is strictly better off under hidden rather than observable testing if and only if preferences lie in the insurance region, i.e. $(\lambda_{DM}, \lambda_A) \in B_{II}$, or $(\lambda_{DM}, \lambda_A) \in W_I$, or $(\lambda_{DM}, \lambda_A) \in W_{II}$.*
2. *The advisor is strictly worse off under hidden rather than observable testing if and only if preferences lie in the skepticism region, i.e. $(\lambda_{DM}, \lambda_A) \in B_I$.*

In the skepticism region, the advisor is strictly worse off under hidden testing, because the DM has a higher acceptance threshold than under observable testing and therefore the advisor is less likely to convince the DM to act in his interest. By contrast, in the insurance region as well as in regions W_I and W_{II} , the advisor is strictly better off under hidden testing, because he can manipulate the DM to act in his interest by strategically withholding information. This illustrates a further point of contrast between the skepticism and the insurance region: in the insurance region, not just the DM but also the advisor strictly benefit from testing being hidden rather than observable.

3.2 N Periods

In this subsection, I show that the key insights of the previous subsection are not specific to a model with $N = 2$ by providing sufficient conditions for their existence for any $N > 2$. For tractability, I will assume that the test is symmetric in the sense that false negatives are equally likely as false positives, i.e. $p_T = p_F = p$.

To build up to comparing the DM's expected payoff under hidden and observable testing, I will first show that under observable testing, it is part of the equilibrium that the advisor stops the flow of information strategically if preferences are sufficiently misaligned. Second, I will show that in the advisor-preferred equilibrium under hidden testing, the advisor pools outcomes to induce the DM to act in his interest.

Given the assumption that the test is symmetric, a sufficient statistic for the posterior belief is how many more positive than negative outcomes were found, independent of the total number of outcomes. I will use x_j to denote the posterior likelihood ratio that the hypothesis is true conditional on observing j more positive than negative outcomes, where $j \in \{-N, \dots, N\}$.

Lemma 2 (Observable Testing) *Suppose testing is observable and the DM prefers to accept for some realizations of the N test outcomes, i.e. $\lambda_{DM} \leq x_N$.*

1. *(Advisor more enthusiastic.) For any $N > 2$, there exists a critical value $\underline{\lambda}_A \leq \lambda_{DM}$, such that if $\lambda_A < \underline{\lambda}_A$ an equilibrium exists in which the following happens on the equilibrium path: The advisor tests initially and stops testing if and only if the likelihood ratio x satisfies $\lambda_{DM} \leq x$. If the advisor stops at $\lambda_{DM} \leq x$ then the DM accepts. If after N periods the likelihood ratio satisfies $x < \lambda_{DM}$ the DM rejects.*
2. *(Advisor more reluctant.) For any $N > 2$, there exists a critical value $\bar{\lambda}_A \geq \lambda_{DM}$, such that if $\lambda_A > \bar{\lambda}_A$, an equilibrium exists in which the following happens on the equilibrium path: the advisor never starts testing and the DM rejects.*
3. *If $\lambda_A < \underline{\lambda}_A$ or if $\lambda_A > \bar{\lambda}_A$ any equilibrium yields the same expected payoff vector.*

Under observable testing, the advisor strategically avoids acquiring evidence if the preferences of the advisor and the DM are sufficiently misaligned. In particular, a sufficiently enthusiastic advisor finds it optimal to stop testing if he has discovered just enough evidence for the DM to accept, while a sufficiently reluctant advisor finds it optimal to never start testing, which guarantees that the DM rejects.

Suppose the advisor is more enthusiastic than the DM. The advisor strictly prefers to keep testing if the evidence collected up to now leads the DM to reject, but for some realization of future outcomes the DM will accept. The reason is that if future outcomes are such that the DM accepts, then the advisor also prefers to accept. Hence, if the outcome is pivotal to the DM's choice, then it changes the DM's choice in line with the advisor's interest. By contrast, the advisor faces a trade-off if the evidence collected up to now leads the DM to accept, but an additional negative outcome would lead her to reject. By testing further, the advisor benefits because he has a chance to find out whether the additional evidence leads both players to agree that rejecting is optimal. However, he risks that the additional evidence leads the DM to reject, but does not convince him that rejecting is optimal. Clearly, the advisor strictly prefers to stop testing in this situation if he is so enthusiastic about accepting that no realization of future outcomes would lead him to reject. A situation in which the advisor always stops as soon as the evidence leads the DM to accept is illustrated in Figure 2.




Figure 2: An example of a combination of preference parameters given $N = 3$ for which the advisor optimally stops testing as soon as the likelihood ratio leads the DM to accept: if the first outcome is positive, the advisor stops testing because the DM acts in line with his interest and the remaining two outcomes can never convince him that rejecting is optimal, not even if both turned out negative.

A reverse reasoning applies if the advisor is more reluctant than the DM. In particular, the advisor faces a trade-off if the evidence collected up to now leads the DM to reject, but an additional positive outcome would lead her to accept. I will find a critical value $\bar{\lambda}_A$ such that an advisor who is more reluctant, i.e. $\bar{\lambda}_A < \lambda_A$, will not test in this situation, although there are realizations of future outcomes for which both players would agree that accepting is optimal. Consider the following example illustrated in Figure 3. Suppose that $N = 3$, the DM accepts if there is at least one more positive outcome than negative outcomes, i.e. $\lambda_{DM} = x_1$, and the advisor accepts if there are at least three more positive than negative outcomes, i.e. $\lambda_A = x_3$. If the first two outcomes are positive ($x = x_2$), then the DM will accept regardless of what the final test outcome will show and, hence, the advisor's expected payoff is independent of this final test. However, at $x = x_2$ the advisor prefers to reject. Therefore, the advisor is strictly better off if he stops after the first test outcome. A second test has no upside, but only the downside that a positive outcome could lead the DM to act against his interest. However, if he stops after a first positive test outcome, then he might as well not start testing. More generally, consider the strongest possible evidence in favor of the hypothesis at which the DM will accept regardless of what the remaining tests show. In the example, this corresponds to the likelihood ratio x_2 . If the advisor prefers to reject at this likelihood ratio, then he is sufficiently reluctant to not start testing.




Figure 3: An example of a combination of preference parameters given $N = 3$ for which the advisor optimally does not start testing, although there exists some realization of test outcomes such that both players agree that accepting is optimal.

Under hidden testing, each player's expected payoff is uniquely defined by restricting attention to advisor-preferred equilibria. The following lemma gives an example of equilibrium play which can generate these expected payoffs. For ease of exposition, let $z \in \{0, \dots, N\}$ denote the total number of positive outcomes obtained in N tests.

Lemma 3 (Hidden Testing) *Suppose testing is hidden.*

1. *An advisor-preferred equilibrium exists and the equilibrium payoff vector is unique.*
2. *(Advisor more enthusiastic.) For any $N > 2$, and for any $\lambda_A < \lambda_{DM}$, the following is an equilibrium path: The advisor runs N tests. There exists a critical level of z , denoted by \bar{z} , such that if $z \geq \bar{z}$ the advisor discloses exactly \bar{z} positive outcomes and the DM accepts, and otherwise he discloses all outcomes and the DM rejects.*
3. *(Advisor more reluctant.) For any $N > 2$, and for any $\lambda_A > \lambda_{DM}$, the following is an equilibrium path: The advisor runs N tests. There exists a critical level of z , denoted by \underline{z} , such that if $z \geq \underline{z}$ the advisor discloses all outcomes and the DM accepts, and otherwise he discloses no outcomes and the DM rejects.*

Under hidden testing, the advisor runs all possible tests. There is an upper bound on the number of positive outcomes disclosed by a more enthusiastic advisor. In equilibrium, when the advisor discloses exactly this number of positive outcomes, the DM infers that the advisor must have found at least this many positive outcomes across all N tests. In an advisor-preferred equilibrium, the advisor chooses the lowest upper bound \bar{z} on the number of positive outcomes disclosed subject to the following two conditions being met. First, the DM optimally accepts conditional on knowing that at least \bar{z} positive outcomes were found in N tests. Second, the advisor prefers to accept conditional on having observed exactly \bar{z} positive outcomes in N tests. By pooling test outcomes in this way, the advisor sometimes achieves that the DM accepts even though the DM would have rejected if she had observed all test outcomes.¹³

By contrast, a more reluctant advisor reveals all outcomes if he prefers to accept conditional on having observed these outcomes and discloses nothing otherwise. The reasoning is the same as for the case with $N = 2$. If the advisor discloses some outcomes, then these outcomes convince the DM to accept. If the advisor discloses nothing, the DM perceives this as bad news and optimally rejects. This shows that a more reluctant advisor can always pool outcomes in a way that allows him to achieve his first-best expected payoff, unlike a more enthusiastic advisor.

The following proposition shows that the skepticism and insurance effects arise for any fixed number of periods $N > 2$.

¹³This pooling of outcomes is similar to the technique of constructing the optimal experiment under Bayesian persuasion in Kamenica and Gentzkow (2011).

Proposition 3 (Skepticism and Insurance Effect)

1. For any $N > 2$, there exists a skepticism region $B_1 \subset \mathbb{R}^2$, i.e. an open set of parameter values $(\lambda_{DM}, \lambda_A)$ in which the DM is strictly better off under hidden than under observable testing, while the advisor is strictly worse off under hidden than under observable testing. In the skepticism region, the advisor is more enthusiastic about accepting than the DM, i.e. $\lambda_A < \lambda_{DM}$.
2. For any $N > 2$, there exists an insurance region $B_2 \subset \mathbb{R}^2$, i.e. an open set of parameter values $(\lambda_{DM}, \lambda_A)$ in which both the DM and the advisor are strictly better off under hidden than under observable testing. In the insurance region, the advisor is more reluctant to accept than the DM, i.e. $\lambda_A > \lambda_{DM}$.

For any $N \geq 2$, the skepticism effect exists for some set of preference parameters. In this set, preferences are sufficiently misaligned such that, under observable testing, the advisor stops testing when the DM is just convinced to accept. This is a disadvantage for the DM, since for some realizations of the remaining outcomes she would have preferred to reject. However, it is possible that under hidden testing, the DM can only be convinced to accept if the advisor shows her sufficiently many positive outcomes such that the DM is willing to accept even if all undisclosed outcomes were negative. In this case, the DM is strictly better off under hidden testing.

For any $N > 2$, the insurance effect also exists for some set of preference parameters. Under observable testing, if the preferences are sufficiently misaligned, the advisor does not start testing and the DM rejects. Under hidden testing, the DM is better off because the advisor reports evidence in favor of the hypothesis if and only if he himself prefers to accept. It is necessary that $N > 2$ if the test is symmetric. Suppose $N = 2$, recall that the DM rejects at the prior belief and suppose the DM accepts if one test is conducted and this test is positive. Then if another test is conducted and it is also positive, the advisor must prefer to accept or he would never accept. However, then the advisor never has a reason to stop after the first test outcome because the DM acts in his interest independent of what the second test shows.

It continues to hold that the DM is never worse off under hidden rather than observable testing if her preferences are sufficiently misaligned with the advisor's preferences.

Proposition 4 (Preference Alignment) *There exists a threshold $d > 0$ such that the DM is weakly better off under hidden rather than observable testing if*

$$|\lambda_A - \lambda_{DM}| > d.$$

As in the case of $N = 2$, there exist parameter combinations for which the DM does not achieve her first-best expected payoff under observable testing, but the skepticism effect allows her to achieve her first-best expected payoff under hidden testing.

Proposition 5 (First-Best Benchmark) *Given any $N > 2$, for preferences in the skepticism region B_I the DM achieves her first-best expected payoff under hidden testing, but not under observable testing.*

4 Delegation and Commitment

This section first compares expected payoffs under delegation to those achieved under information acquisition and then studies how the comparison between hidden and observable testing changes when either party has commitment power. This analysis helps to further illustrate differences between the skepticism and the insurance effect. Throughout I allow for $N \geq 2$ and assume that the test has accuracy p independent of the state.

The following proposition shows that the DM weakly prefers information transmission to delegation if testing is hidden, regardless of how aligned preferences are. However, the same is not true if testing is observable.

Proposition 6 (Delegation) *Suppose the DM could delegate decision rights to the advisor.*

1. *The DM is never strictly better off under delegation than hidden testing.*
2. *For any $N \geq 2$, there exists an open set $(\lambda_{DM}, \lambda_A) \subset \mathbb{R}^2$ for which the DM is strictly better off under delegation than observable testing.*
3. *For $N = 2$ then the DM is strictly better off under delegation than observable testing if and only if preferences lie in the insurance region B_{II} .*

Under delegation, the advisor runs all tests and accepts if and only if he finds it optimal to accept. A disadvantage of delegation for the DM is that, for some test outcomes and some preferences of the advisor, the advisor chooses a different action than what the DM finds optimal. However, an advantage of delegation can exist if the advisor has incentives to conduct more tests than he would without delegation.

Under hidden testing, if the outcomes are such that the players agree on what the optimal action is then this action is implemented, because otherwise the advisor would have a profitable deviation to reveal all outcomes. The same is true under delegation. However, if outcomes are such that the players disagree, then under delegation, the advisor's preferred choice is implemented for a weakly larger set of outcome realizations than the DM's preferred choice. The DM does not benefit from giving decision rights to the advisor if he has already given him the discretion to selectively disclose outcomes.

However, when testing is observable, then delegation can increase the advisor's incentives to test. This is because delegation eliminates the risk that additional outcomes lead the DM to

act against the advisor's interest, i.e. it generates the insurance effect. If $N = 2$ then this is beneficial for the DM if and only if parameters lie in the insurance region B_{II} . With $N > 2$ periods, delegation is not necessarily beneficial only in the insurance region. Without delegation it is possible that the advisor stops testing once the evidence is just strong enough for the DM to accept, even if it is possible that the remaining outcomes could provide such strong evidence against the hypothesis being true that both players would agree to reject. The reason the advisor stops testing is that it is too likely that the remaining outcomes provide only weak evidence against the hypothesis being true, causing the DM to reject when the advisor would prefer to accept.¹⁴ In this situation, i.e. conditional on the advisor stopping testing, delegation is beneficial for the DM. If the decision was delegated to the advisor, he would continue testing because he can ensure that rejection is chosen if and only if he prefers rejection given the evidence collected.¹⁵

So far I have assumed that the DM has no commitment power. In what follows, I will explore what would change if the DM could commit ex ante to which action she will take for any evidence presented to her.

Lemma 4 (DM Commitment) *If the DM has the power to commit ex ante to her actions contingent on outcomes, then she achieves her first-best expected payoff, whether testing is hidden or observable.*

If the advisor is more enthusiastic about accepting than the DM, the DM can achieve her first-best expected payoff by committing to reject if the advisor presents fewer than N outcomes to her. Similarly, if the advisor is more reluctant to accept than the DM, the DM can achieve her first-best expected payoff by committing to accept if the advisor presents fewer than N outcomes to her. To compare the DM's welfare across different settings, denote her expected payoff with commitment power by $\pi_{DM}(C)$ and denote her expected payoff without commitment by $\pi_{DM}(NC, i)$ where $i \in \{OT, HT\}$ indicates whether testing is observable (OT) or hidden (HT).

Proposition 7 (DM Commitment) *Suppose testing is observable and the DM does not have the power to commit to actions contingent on outcomes.*

1. *For preferences in the skepticism region B_I , the DM's marginal benefit from allowing testing to be hidden is equally high as her marginal benefit from gaining commitment power:*

$$\pi_{DM}(NC, OT) < \pi_{DM}(NC, HT) = \pi_{DM}(C).$$

¹⁴This effect can be seen as a mirror image of the insurance effect.

¹⁵However, from an ex-ante perspective, there is also a downside to delegation. There are some draws of test outcomes for which, under observable testing, the advisor does not stop testing and the DM optimally rejects, whereas under delegation the advisor accepts.

2. For preferences in the insurance region B_{II} , the DM's marginal benefit from allowing testing to be hidden is smaller than her marginal benefit from gaining commitment power:

$$\pi_{DM}(NC, OT) < \pi_{DM}(NC, HT) < \pi_{DM}(C).$$

This proposition shows that, in some circumstances, the DM can use hidden testing to circumvent a commitment problem. In particular, both hidden testing and commitment can be used to generate the skepticism effect if $(\lambda_{DM}, \lambda_A) \in B_1$. For example, if $N = 2$, the DM accepts if she knows one test was run and the outcome was positive, but ideally she would like to make her decision dependent on whether or not an additional test also shows a positive outcome. To generate the skepticism effect under observable testing, the DM needs to commit to act less in the advisor's interest. In particular, she needs to commit to reject based on a single positive outcome. By contrast, under hidden testing, the threat to reject a single positive outcome is credible even without commitment power, because the DM has reason to suspect that the advisor has omitted some negative evidence.

Commitment can be used to generate the insurance effect, but it allows the DM to do even better. Suppose preferences lie in the insurance region B_2 . Under observable testing, the DM can induce the advisor to test by committing to act in his interest for any realization of outcomes. Under hidden testing, such commitment is not necessary because the advisor has the possibility to omit outcomes and thereby prevent the DM from accepting if he would like to reject. However, the DM could do even better if she commits to accept unless the advisor shows her evidence that leads her to reject, i.e. she commits to act less in the advisor's interest. With this commitment the DM can achieve her first-best expected payoff.

It is also insightful to see what would happen if the advisor had the power to commit ex ante to how many tests he runs and to which outcomes he discloses given his findings.

Proposition 8 (Advisor Commitment) *Suppose the advisor has the power to commit to a testing and disclosure strategy ex ante.*

1. *If the advisor is more enthusiastic about accepting than the DM, then the DM is never strictly better off under hidden rather than observable testing.*
2. *If the advisor is more reluctant to accept than the DM, then the DM is strictly better off under hidden testing than observable testing if preferences lie in the insurance region B_2 .*

Suppose the advisor is more enthusiastic about accepting than the DM. If $N = 2$ and the advisor has no commitment power, the DM benefits from hidden testing for preferences in the skepticism region B_I , because she can credibly reject based on the report of a single positive outcome. By contrast, if the advisor has commitment power then he would optimally commit to reporting a

single positive outcome if and only if this outcome was obtained on the first test. As a consequence, it is no longer possible for the DM to credibly reject a single positive outcome, just as in the case of observable testing. Put differently, when the advisor has commitment power he can do at least as well as under observable testing and therefore the skepticism effect ceases to exist. If the advisor is more reluctant to accept than the DM then he achieves the same expected payoff under hidden testing as when the decision was delegated to him. Therefore, he does not strictly benefit from commitment. In this situation, the DM's payoff comparison between hidden and observable testing is the same as when the advisor has no commitment power.

5 Robustness Checks

This section revisits the key comparison between expected payoffs under hidden and observable testing when i) the advisor has to commit to the number of tests in advance, i.e. testing is simultaneous rather than sequential and ii) the horizon is infinite and the advisor incurs a constant cost per test. I will again assume the test has accuracy p independent of the state.

In the preceding analysis, I assumed that the advisor acquires test outcomes sequentially. In some circumstances, the advisor may have to decide in advance how much evidence to acquire before he learns about any of the outcomes. If $N = 2$, the results in Theorem 1 still apply. However, if $N > 2$ then there is not always a region of preference parameters for which the skepticism effect exists.

Proposition 9 (Simultaneous Testing) *Suppose $N > 2$ and the advisor has to commit to the number of tests ex ante.*

1. *The skepticism effect does not exist for all N , i.e. for some N there does not exist an open set $(\lambda_{DM}, \lambda_A) \subset \mathbb{R}^2$ such that the DM is strictly better off under hidden rather than observable testing.*
2. *The insurance effect exists for all N , i.e. for any N , there exists some open set $(\lambda_{DM}, \lambda_A) \subset \mathbb{R}^2$ such that the DM is strictly better off under hidden rather than observable testing.*

Suppose the advisor is more enthusiastic. When testing is sequential, the advisor has the option to stop testing conditional on having acquired enough evidence to convince the DM to accept. When testing is simultaneous, however, the advisor does not know whether or not the evidence would be sufficient to convince the DM to accept. In expectation, it is more likely that further test outcomes lead the DM to accept when he would have otherwise rejected rather than lead the DM to reject when he would have otherwise accepted. Therefore, it is possible that the advisor runs all N tests for any configuration of preference parameters under observable testing and, hence, the DM cannot be strictly better off under hidden testing. This reasoning does not apply if the

advisor is more reluctant. If preferences are sufficiently misaligned it is optimal not to test at all, whether testing is simultaneous or sequential. This is because a sufficient number of tests needs to be conducted in order for the advisor to ever be convinced to accept. But if this number of tests is conducted, the drawback that with some chance the DM accepts when the advisor wants to reject can outweigh the benefit that the DM accepts when the advisor wants to accept.

Furthermore, the preceding analysis was built on the assumption that the horizon is finite and tests are costless. This finite horizon was used to represent the fact that the advisor is constrained in the total resources he can devote to testing. In some circumstances, the advisor may not face such constraints, i.e. the horizon may be infinite. Given an infinite horizon and costless tests, analyzing the conflict of interest between the advisor and the DM would cease to be interesting, since they agree on the optimal action under certainty. To avoid this, I introduce a constant cost per test.¹⁶

Proposition 10 (Infinite Horizon) *Suppose $N \rightarrow \infty$ and the advisor incurs a constant cost per test.*

1. *There exist parameter combinations for which the skepticism effect exists.*
2. *There exists no parameter combination for which the insurance effect exists.*

The reason for why the skepticism effect exists is very similar to the one discussed in the preceding analysis with a finite horizon and costless testing. Suppose $\lambda_A = 0$, i.e. the advisor weakly prefers the DM to accept independent of the state. Then under observable testing, the advisor stops testing for one of two reasons. Either he has found enough evidence to convince the DM to accept or he gives up because he expects that it would be too costly to convince the DM to accept given his evidence. Therefore, on the equilibrium path, the DM accepts if and only if the evidence is just strong enough to lead him to accept. There exist parameter combinations for which the DM requires a larger number of positive outcomes to accept when testing is hidden. As in the preceding analysis, the reason is that the advisor can hide negative outcomes and, therefore, there are histories of test outcomes at which he can stop and convince the DM to accept even if under observable testing the DM would have rejected. As the DM's threshold of acceptance rises, it becomes relatively more likely that the advisor gives up when the hypothesis is false relative to when it is true and, therefore, his report becomes more informative.¹⁷ This benefits the DM.

However, an insurance effect does not exist in this alternative model. To understand why, the first step is to realize that the advisor's optimal testing strategy is stationary and depends only on the current posterior belief. Under observable testing, this follows from the fact the DM's optimal action choice depends only on the posterior belief. Under hidden testing, this follows because in

¹⁶Felgenhauer and Schulte (2014) study an infinite-horizon model of hidden testing with a constant cost of testing and an advisor with state-independent preferences.

¹⁷Felgenhauer and Schulte (2014) show that an increase in the DM's threshold for acceptance improves how informative it is that the advisor meets the threshold.

equilibrium the more reluctant advisor's reporting strategy leads the DM to act in his interest at any posterior belief, as was the case in the preceding analysis. A necessary condition for the insurance effect to exist is that, under hidden testing, the advisor must find it worthwhile to keep testing at beliefs which lie between the prior belief and the belief at which he just prefers acceptance. Otherwise, the advisor will never find evidence that leads him to prefer acceptance. However, if the advisor optimally keeps testing in this range of beliefs under hidden testing, then he must also optimally keep testing under observable testing. This is because under observable testing stopping is an even less attractive option for the advisor than under hidden testing, as it can result in the DM acting against his interest. Therefore, the insurance that the DM acts in the advisor's interest under hidden testing does not lead to additional information acquisition.

For the insurance effect to exist, the cost per test must be increasing in the number of tests, e.g. a finite horizon can be interpreted as an extreme case of a convex cost of testing. The reason is that when costs are convex, the advisor's optimal testing strategy depends not only on his posterior belief, but also on how much time has passed. Therefore, it might be that the posterior belief is such that the DM accepts when the advisor would prefer to reject, yet the advisor stops testing because the cost has become too high (or he has run out of time). To avoid finding himself in such a situation, the advisor may instead stop testing sooner at a posterior belief at which both players agree that rejecting is optimal. By contrast, under hidden testing, since the DM always acts in the advisor's interest, the consequences of stopping at certain histories of outcomes are not as negative as under observable testing and, hence, the advisor may acquire more information than under observable testing.

In many situations, it is reasonable to assume that the cost of testing is convex, e.g. a pharmaceutical company may find it increasingly difficult to recruit subjects for their trials the more trials they run, or that overall resources are limited, e.g. the budget for development of a drug is limited. Similar consequences would arise if with some probability the final decision as to be taken in a given period and this probability increases in the number of periods.

6 Related Literature

My work builds on the extensive literature on persuasion with verifiable information.¹⁸ Like many papers in this literature, I assume that information acquisition is endogenous.¹⁹ In particular, my work belongs to the part of the literature studying situations in which one player is faced with a choice under uncertainty, but relies on another player to acquire information and no contracts can be written.²⁰ The papers most closely related to mine are those which compare the DM's

¹⁸E.g. Grossman (1981), Milgrom (1981), Milgrom and Roberts (1986), Shin (1994)

¹⁹E.g. Farrell and Sobel (1983), Matthews and Postlewaite (1985), Shavell (1994)

²⁰E.g. Brocas and Carrillo (2007), Che and Kartik (2009), Kamenica and Gentzkow (2011)

payoff under hidden and observable information acquisition, e.g. Brocas and Carrillo (2007), Henry (2009) and Felgenhauer and Loerke (2017), which will be discussed in more detail below. In addition, Henry and Ottaviani (2017) study welfare effects and the impact of commitment and delegation in a model of observable sequential information acquisition. Focusing exclusively on hidden information acquisition, Felgenhauer and Schulte (2014) study how variations in the cost of experimentation affect persuasion.

My work is related to the literature on the ideal bias of an advisor when the DM can choose how closely aligned the advisor's preferences are with her own, e.g. Che and Kartik (2009), Gerardi and Yariv (2008) and Dur and Swank (2005). For a more detailed comparison with Che and Kartik (2009) see below. In contrast to this literature, I analyze how the benefits of hidden testing vary with the advisor's bias. Second, my work contributes to the literature on delegation of decision rights.²¹ In contrast to Li and Suen (2004) who use the same assumptions on preferences but take information to be exogenously given, I show that there are situations in which the DM is worse off by delegating the decision-making to a more reluctant advisor. In a cheap talk setting, Argenziano et al. (2016) and Deimen and Szalay (2015) show that the DM prefers communication to delegation when the advisor acquires information. Assuming verifiable disclosure, I show that the DM prefers information transmission to delegation irrespective of the advisor's bias if testing is hidden, but that this finding does not hold if testing is observable. That more discretion can have ambiguous effects on the DM's payoff is also illustrated by Di Tillio et al. (2017). They show that for some distributions the DM is better off if she allows a biased advisor to collect a sample of size n in private and report his preferred observation than if she restricts him to only collect a single observation. Third, my work is related to the literature on strategic commitment.²² When the DM cannot commit to how she takes decisions based on outcomes, she may benefit from giving up the option to observe outcomes directly.

I will now turn to discuss the comparison to the most related papers in detail. Brocas and Carrillo (2007) were the first to study a persuasion game in which the advisor sequentially acquires binary signals. They assume that the DM has three options and the advisor's preferences over these options are state-independent. Under observable testing, they show that it is optimal for the advisor to stop if the DM choose his most preferred option given the evidence. This result also applies in my setting when the advisor has state-independent preferences, i.e. when $\lambda_A = 0$. Brocas and Carrillo (2007) extend the model to allow for information acquisition being costly and also consider a variant of the model in which the DM has the possibility to acquire information. They find that the players' payoffs are unchanged under hidden testing when the advisor can verifiably reveal his posterior belief. In contrast to their work, I assume that the advisor can verifiably disclose outcomes, but not his posterior beliefs. To illustrate this difference, suppose the advisor has

²¹E.g. see Aghion and Tirole (1997), Dessein (2002) and Li and Suen (2004) when information is exogenous.

²²E.g. Schelling (1960)

observed one negative and then one positive outcome. In my setting, he can report a single positive outcome, although his posterior belief is not equal to the posterior belief he would have if he had only observed one positive outcome. With my assumption on verifiable disclosure, there exist equilibria in which payoffs differ depending whether testing is hidden or observable. In particular, if the advisor-preferred equilibrium is selected, the DM may strictly benefit from hidden testing.

Henry (2009) finds that the DM is always weakly better off with hidden than with observable testing, and sometimes strictly better off. In his setting, the state and action space are continuous, the players' loss functions are quadratic and the advisor's ideal action differs from the DM's ideal action by a constant bias independent of the state. The advisor commits ex ante to a quantity of research, which maps into a number of infinitesimal positive and negative signals drawn from a state-dependent distribution. Increasing the quantity of research is costly. Henry (2009) shows that under hidden testing, the advisor acquires more signals for any given response by the DM, because he benefits from having a larger range of signals to choose from. In equilibrium, the DM can perfectly infer how much research the advisor has conducted. Since any incremental change in the DM's posterior belief about the state affects her action choice, Milgrom and Roberts (1986)'s unraveling argument applies. In any sequential equilibrium, the DM optimally interprets any withheld signal as unfavorable evidence, i.e. he adopts a skeptical attitude, and the advisor has an incentive to reveal all favorable evidence. Henry (2009) also analyzes the contrast between observable and hidden testing when the DM is uncertain about the advisor's bias and when two competing advisors make recommendations. In my setting, the unraveling argument does not apply because the DM's choice is discrete.²³ When the advisor has to commit ex ante to how many signals he collects, there is not always a combination of preference parameters for which the DM is strictly better off under hidden testing (see Proposition 9). Even if the advisor does not commit ex ante to how many signals he collects, there exist equilibria in which the advisor can omit evidence without affecting which action the DM chooses. Consequently, the DM may be worse off under hidden testing. In addition, it is worth noting that Henry (2009) allows for the advisor's ideal action to be biased upwards or downwards relative to the DM's ideal action, yet the insurance effect never arises. This is because in Henry (2009)'s setting, for any belief, the DM and the advisor disagree on what the optimal action is and, therefore, the advisor always has a reason to persuade the DM.

Felgenhauer and Loerke (2017) also show that the DM is always weakly better off under hidden than under observable testing. In their setting, the advisor cannot only decide how many tests to run, but can also design how informative each test will be. When the advisor discloses an outcome he also discloses the technology of the test by which it was generated. State and action spaces are binary, tests are costly, the advisor prefers to accept independent of the state (i.e. $\lambda_A = 0$), but the DM has a loss from taking an inappropriate action. Interestingly, they find that the advisor

²³Therefore, the DM's expected payoff function is not continuously differentiable and strictly concave in her choice variable. As a consequence, changes in posterior beliefs do not necessarily translate into changes in optimal choices.

runs only a single test in equilibrium, whether testing is observable or hidden. The DM benefits from hidden testing because the advisor chooses a more informative test and discloses fully. By running a more informative test, the advisor can credibly commit that he will not run any further tests even if the outcome is negative. This is because given his posterior belief following a negative outcome, the expected cost of convincing the DM outweighs the expected benefit. In my setting, if $\lambda_A = 0$, hidden testing can be beneficial because the DM's skepticism induces the advisor to run additional tests so that he can show that these back up earlier findings, but hidden testing can also be detrimental if the DM accepts despite anticipating that the advisor may have omitted negative outcomes. In many real-life situations, an advisor would have some flexibility with regard to the test he runs, but not full flexibility. Felgenhauer and Loerke (2017) study one extreme, while I focus on the other extreme. Comparing our findings illustrates that whether or not hidden testing is beneficial for the DM depends on what assumptions are made about how much flexibility the advisor has when designing tests.

Che and Kartik (2009) primarily study a persuasion game when preferences are aligned, but the players hold different prior beliefs. More closely related to my work is the part of their paper where they assume differences in preferences but not in prior beliefs. They analyze how the information obtained by the DM varies as the two players' preferences become less aligned, when testing is hidden or when it is observable. In their setting, the state and action spaces are continuous, the players' loss functions are quadratic and the advisor's ideal action differs from the DM's ideal action by a constant bias independent of the state. The advisor exerts costly effort to increase the chances to observe a single signal which is normally distributed about the true state. They show that under observable testing, an increase in misalignment of preferences does not affect the DM's expected payoff. However, under hidden testing, a larger conflict of interest may leave the DM better off. A more biased advisor does not report the signal for a larger range of realizations. Since the DM anticipates this, she optimally responds less favorably if no signal is reported. Hence, the more biased advisor has more incentives to exert effort. In my setting, the advisor's bias affects information acquisition even if testing is observable, due to the assumption that players' interests are aligned under certainty. In addition, if the comparison between hidden and observable testing was made in their setting, then it would show that for some preference combinations, hidden testing can be detrimental for the DM. This is because the advisor's incentive to exert effort do not increase enough to compensate for the fact that he can strategically disclose signals. By contrast, in my setting, hidden testing can be detrimental because the advisor runs additional tests which are not valuable to the DM, but help the advisor to manipulate his disclosure. Finally, whereas Che and Kartik (2009) find that delegating the decision making to a more biased advisor makes the DM worse off, I find that delegating to a more reluctant advisor can make the DM better off.

7 Conclusion

When taking decisions under uncertainty we often have to rely on others to carry out research. Since their preferences are not always aligned with ours, a natural question is whether or not we should monitor more closely how they go about collecting information. This paper shows that making the collection of information transparent can have adverse consequences for decision making. On the one hand, transparency ensures that an advisor cannot strategically omit findings. On the other hand, transparency can discourage information acquisition, even if the advisor and the decision maker agree on the optimal action under certainty.

The paper has distinguished between two effects which cause the DM to be strictly better off when information acquisition is private: the skepticism effect, which can arise when the advisor is more enthusiastic about accepting than the DM, and the insurance effect, which can arise when the advisor is more reluctant about accepting than the DM. The skepticism effect shows that when information acquisition is private the DM can credibly raise her threshold for acceptance because she has reason to suspect that the advisor is hiding contradicting evidence. By contrast, the insurance effect shows that when information acquisition is private the advisor does not face the risk that additional evidence leads the DM to act against his interest. Therefore, he can freely explore whether or not additional evidence is sufficiently strong for both to agree on a different action choice. Due to the presence of these effects, there are combination of preference parameters for which the DM achieves her first-best expected payoff under private, but not under public information acquisition. While the insurance effect also causes the advisor's expected payoff to increase, the skepticism effect causes the advisor's expected payoff to decrease.

In addition, this paper has shown that there are combination of preferences for which the DM is strictly worse off when information acquisition is private. A sufficient condition for the DM to be weakly better off under private information acquisition is for preference to be sufficiently misaligned, i.e. for differences in the relative loss from inappropriate actions to be large between players.

While the current paper analyzes the interaction between an individual DM and advisor, there are other interesting consequences of transparency when there are several decision makers using the evidence as a basis for their choice or several advisors supplying evidence. I leave this for future research.

A Appendix

A.1 Lemma 5 [Observable Testing Equilibrium]

To characterize the players' expected losses in equilibrium under observable testing, I will use the following methodology. Suppose Nature draws the outcomes of both tests before the advisor moves. Denote the joint draws by

$$v \in V = \{(+, +), (+, -), (-, +), (-, -)\}.$$

The draws v are unobserved by either player, but if the advisor runs a test in period 1 (2), he observes the first (second) element of v . Any equilibrium can be represented as a function e from V into the set of DM's actions, i.e.

$$e : V \rightarrow T = \{accept, reject\}.$$

If any two equilibria are represented by the same function e then these equilibria yield the same expected loss for each player. Lemma 5 defines this equilibrium function for the game with observable testing.

Lemma 5 (Observable Testing on Path)

1. In Region 1, i.e. if

$$x_{\{\emptyset\}} < \lambda_{DM} \leq x_{\{+,-\}},$$

then

- if $\lambda_A \leq x_{\{+,-\}}$ the DM accepts iff $v \in \{(+, +), (+, -), (-, +)\}$, and
- if $x_{\{+,-\}} < \lambda_A \leq x_{\{+\}}$ the DM accepts iff $v \in \{(+, +), (+, -)\}$, and
- if $\lambda_A > x_{\{+\}}$ the DM never accepts.

2. In Region 2, i.e. if

$$\max \{x_{\{\emptyset\}}, x_{\{+,-\}}\} < \lambda_{DM} \leq x_{\{+\}},$$

then

- if $\lambda_A \leq x_{\{+,-\}}$ the DM accepts iff $v \in \{(+, +), (+, -)\}$, and
- if $x_{\{+,-\}} < \lambda_A < x_{\{+,+\}}$ the DM accepts iff $v \in \{(+, +)\}$, and
- if $\lambda_A > x_{\{+,+\}}$ the DM never accepts.




Figure 4: This Figure indicates along the vertical axis the regions in which the DM's optimal decision rule under observable testing is unchanged, in $(\lambda_A, \lambda_{DM})$ –space given $p_T > p_F$.

3. In Region 3, i.e. if

$$x_{\{+\}} < \lambda_{DM} \leq x_{\{+,+\}},$$

then

- if $\lambda_A \leq x_{\{+,+\}}$ the DM accepts iff $v \in \{(+, +)\}$, and
- if $\lambda_A > x_{\{+,+\}}$ the DM never accepts.

4. In Region 4, i.e. if

$$\lambda_{DM} > x_{\{+,+\}},$$

then the DM never accepts.

Proof of Lemma 5: As described in Section 3.1, the DM optimally accepts if and only if her posterior likelihood ratio exceeds λ_{DM} . Due to the discreteness of the set-up, it is helpful to divide the range of $\lambda_{DM} \in (x_{\{\emptyset\}}, \infty)$ into four regions such that the DM's optimal strategy is unchanged within each region. These regions are indicated along the vertical axis in Figure 4. Region 1 is empty unless $p_T > p_F$, since $x_{\{\emptyset\}} < x_{\{+,-\}}$ if and only if $p_T > p_F$. For each region, I will derive the advisor's optimal strategy using backward induction and, based on this, find the subset of draws of outcomes for which the DM accepts. Throughout, denote the advisor's strategy by

$$\sigma_A \equiv (\sigma_A(h_0), \sigma_A(h_1 = (+)), \sigma_A(h_1 = (-)), \sigma_A(h_1 = (\emptyset))) \in [0, 1]^4.$$

Region 1. Suppose $p_T > p_F$ and $x_{\{\emptyset\}} < \lambda_{DM} \leq x_{\{+,-\}}$.

Equilibrium strategies are as follows: The DM accepts if and only if h_2 is such that $\tilde{h} \in \{\{+, +\}, \{+\}, \{+, -\}\}$. For $\lambda_A \leq x_{\{+,-\}}$, the advisor's strategy is $\sigma_A = (1 - \varepsilon, \varepsilon, 1 - \varepsilon, 1 - \varepsilon)$. For $x_{\{+,-\}} < \lambda_A \leq x_{\{+\}}$, the advisor's strategy is $\sigma_A = (1 - \varepsilon, \varepsilon, \varepsilon, 1 - \varepsilon)$. For $x_{\{+\}} < \lambda_A$, the advisor's strategy is $\sigma_A = (\varepsilon, \varepsilon, \varepsilon, \varepsilon)$. Beliefs at h_n are given by the posterior likelihood ratio $\frac{Pr(true|h_n)}{Pr(false|h_n)}$.

Suppose $\varepsilon \rightarrow 0$. At the end of period 2, the DM's strategy is optimal since

$$x_{\{-,-\}} < x_{\{-\}} < x_{\{\emptyset\}} < \lambda_{DM} \leq x_{\{+,-\}} < x_{\{+\}} < x_{\{+,+\}}.$$

At the start of period 2, given $h_1 = (+)$, the DM accepts independent of whether or not another test is run. Hence, the advisor is indifferent between running another test or not. Given $h_1 = (-)$, if the advisor does not test the DM rejects. If the advisor tests then the DM accepts if and only if the outcome is positive. The advisor strictly prefers not to test if and only if

$$\begin{aligned} Pr(true|h_1 = (-)) &< \lambda_A Pr(false|h_1 = (-))(1 - p_F) + Pr(true|h_1 = (-))(1 - p_T) \\ \lambda_A &> \frac{Pr(true|h_1 = (-)) p_T}{Pr(false|h_1 = (-))(1 - p_F)} = \frac{q p_T (1 - p_T)}{(1 - q)(1 - p_F) p_F} \equiv x_{\{+,-\}}, \end{aligned}$$

since $\frac{Pr(true|h_1 = (-))}{Pr(false|h_1 = (-))} = \frac{q(1 - p_T)}{(1 - q)p_F}$. Otherwise, he strictly prefers to test. Given $h_1 = (\emptyset)$, if the advisor does not test the DM rejects. If the advisor tests then the DM accepts if and only if the outcome is positive. The advisor strictly prefers not to test if and only if

$$\begin{aligned} Pr(true) &< \lambda_A Pr(false)(1 - p_F) + Pr(true)(1 - p_T) \\ \lambda_A &> \frac{Pr(true) p_T}{Pr(false)(1 - p_F)} = \frac{q p_T}{(1 - q)(1 - p_F)} \equiv x_{\{+\}}. \end{aligned}$$

Otherwise, he strictly prefers to test.

At the start of period 1, if the advisor does not test the DM rejects. For $\lambda_A \leq x_{\{+,-\}}$, if he tests then he will run another test if the first outcome is negative. He strictly prefers to test since

$$\begin{aligned} Pr(true) &> \lambda_A Pr(false)(1 - p_F^2) + Pr(true)(1 - p_T)^2 \\ \lambda_A &< \frac{q(1 - (1 - p_T)^2)}{(1 - q)(1 - p_F^2)} \end{aligned}$$

and

$$\frac{q(1 - (1 - p_T)^2)}{(1 - q)(1 - p_F^2)} < x_{\{+,-\}}.$$

For $\lambda_A > x_{\{+,-\}}$, if he tests then he will stop testing if the first outcome is negative. The DM will accept if and only if the first outcome is positive. The advisor strictly prefers not to test if and only

if

$$\begin{aligned} Pr(true) &< \lambda_A Pr(false)(1 - p_F) + Pr(true)(1 - p_T) \\ \lambda_A &> \frac{qp_T}{(1 - q)(1 - p_F)} \equiv x_{\{+\}}. \end{aligned}$$

Otherwise, he strictly prefers to test.

Region 2. Suppose $\max\{x_{\{\emptyset\}}, x_{\{+,-\}}\} < \lambda_{DM} \leq x_{\{+\}}.$

Equilibrium strategies are as follows: The DM accepts if and only if h_2 is such that $\tilde{h} \in \{\{+, +\}, \{+\}\}.$ For $\lambda_A \leq x_{\{+,-\}},$ the advisor's strategy is $\sigma_A = (1 - \varepsilon, \varepsilon, \varepsilon, 1 - \varepsilon).$ For $x_{\{+,-\}} < \lambda_A \leq x_{\{+\}},$ the advisor's strategy is $\sigma_A = (1 - \varepsilon, 1 - \varepsilon, \varepsilon, 1 - \varepsilon).$ For $x_{\{+\}} < \lambda_A \leq x_{\{+,+\}},$ the advisor's strategy is $\sigma_A = (1 - \varepsilon, 1 - \varepsilon, \varepsilon, \varepsilon).$ For $x_{\{+,+\}} < \lambda_A,$ the advisor's strategy is $\sigma_A = (\varepsilon, 1 - \varepsilon, \varepsilon, \varepsilon).$ Beliefs at h_n are given by the posterior likelihood ratio $\frac{Pr(true|h_n)}{Pr(false|h_n)}.$

Suppose $\varepsilon \rightarrow 0.$ At the end of period 2, the DM accepts if and only if $\tilde{h} \in \{\{+, +\}, \{+\}\},$ since

$$x_{\{-,-\}} < x_{\{-\}} < \min\{x_{\{\emptyset\}}, x_{\{+,-\}}\} < \max\{x_{\{\emptyset\}}, x_{\{+,-\}}\} < \lambda_{DM} \leq x_{\{+\}} < x_{\{+,+\}}.$$

At the start of period 2, given $h_1 = (-),$ the DM rejects independent of whether or not another test is run. Hence, the advisor is indifferent between running another test or not. Given $h_1 = (+),$ if the advisor does not test the DM accepts. If the advisor tests then the DM accepts if and only if the outcome is positive. The advisor strictly prefers not to test if and only if

$$\begin{aligned} \lambda_A Pr(false|h_1 = (+)) &< \lambda_A Pr(false|h_1 = (+))(1 - p_F) + Pr(true|h_1 = (+))(1 - p_T) \\ \lambda_A &< \frac{Pr(true|h_1 = (+))(1 - p_T)}{Pr(false|h_1 = (+))p_F} = \frac{qp_T(1 - p_T)}{(1 - q)(1 - p_F)p_F} \equiv x_{\{+,-\}} \end{aligned}$$

since $\frac{Pr(true|h_1 = (+))}{Pr(false|h_1 = (+))} = \frac{qp_T}{(1 - q)(1 - p_F)}.$ Otherwise, he strictly prefers to test. Given $h_1 = (\emptyset),$ if the advisor does not test the DM rejects. If the advisor tests then the DM accepts if and only if the outcome is positive. The advisor strictly prefers not to test if and only if

$$\begin{aligned} Pr(true) &< \lambda_A Pr(false)(1 - p_F) + Pr(true)(1 - p_T) \\ \lambda_A &> \frac{qp_T}{(1 - q)(1 - p_F)} \equiv x_{\{+\}}. \end{aligned}$$

Otherwise, he strictly prefers to test.

At the start of period 1, if the advisor does not test the DM rejects. For $\lambda_A \leq x_{\{+,-\}},$ if the advisor tests then he will stop testing if the first outcome is positive. The DM will accept if and

only if the first outcome is positive. It is strictly optimal for the advisor to test since

$$\begin{aligned} Pr(true) &> \lambda_A Pr(false)(1 - p_F) + Pr(true)(1 - p_T) \\ \lambda_A &< \frac{qp_T}{(1 - q)(1 - p_F)} \equiv x_{\{+\}}. \end{aligned}$$

For $\lambda_A > x_{\{+,-\}}$, if the advisor tests then he will not stop testing if the first outcome is positive. The DM will accept if and only if both outcomes are positive. The advisor strictly prefers not to test if and only if

$$\begin{aligned} Pr(true) &< \lambda_A Pr(false)(1 - p_F)^2 + Pr(true)(1 - p_T^2) \\ \lambda_A &> \frac{qp_T^2}{(1 - q)(1 - p_F)^2} \equiv x_{\{+,+\}}. \end{aligned}$$

Otherwise, he strictly prefers to test.

Region 3. Suppose $x_{\{+\}} < \lambda_{DM} \leq x_{\{+,+\}}$.

Equilibrium strategies are as follows: The DM accepts if and only if h_2 is such that $\tilde{h} \in \{\{+,+\}\}$. For $\lambda_A \leq x_{\{+,+\}}$, the advisor's strategy is $\sigma_A = (1 - \varepsilon, 1 - \varepsilon, \varepsilon, 1 - \varepsilon)$. For $x_{\{+,+\}} < \lambda_A$, the advisor's strategy is $\sigma_A = (\varepsilon, \varepsilon, \varepsilon, \varepsilon)$. Beliefs at h_n are given by the posterior likelihood ratio $\frac{Pr(true|h_n)}{Pr(false|h_n)}$.

Suppose $\varepsilon \rightarrow 0$. At the end of period 2, the DM accepts if and only if $\tilde{h} \in \{\{+,+\}\}$, since

$$x_{\{-,-\}} < x_{\{-\}} < \min\{x_{\{\emptyset\}}, x_{\{+,-\}}\} < \max\{x_{\{\emptyset\}}, x_{\{+,-\}}\} < x_{\{+\}} < \lambda_{DM} \leq x_{\{+,+\}}.$$

At the start of period 2, given $h_1 = (-)$, the DM rejects independent of whether or not another test is run. Hence, the advisor is indifferent between running another test or not. Given $h_1 = (+)$, if the advisor does not test the DM rejects. If the advisor tests then the DM accepts if and only if the outcome is positive. The advisor strictly prefers not to test if and only if

$$\begin{aligned} Pr(true|h_1 = (+)) &< \lambda_A Pr(false|h_1 = (+))(1 - p_F) + Pr(true|h_1 = (+))(1 - p_T) \\ \lambda_A &> \frac{Pr(true|h_1 = (+))p_T}{Pr(false|h_1 = (+))(1 - p_F)} = \frac{qp_T^2}{(1 - q)(1 - p_F)^2} \equiv x_{\{+,+\}}, \end{aligned}$$

since $\frac{Pr(true|h_1 = (+))}{Pr(false|h_1 = (+))} = \frac{qp_T}{(1 - q)(1 - p_F)}$. Otherwise, he strictly prefers to test. Given $h_1 = (\emptyset)$, the DM rejects independent of whether another test is run or not. The advisor is indifferent between testing or not.

At the start of period 1, if the advisor does not test the DM rejects. For $\lambda_A \leq x_{\{+,+\}}$, since the advisor will test at $h_1 = (+)$ he prefers to test. For $\lambda_A > x_{\{+,+\}}$, the advisor will stop testing after a first positive outcome. The DM will therefore always reject independent of whether another test

is run or not. The advisor is indifferent between testing or not.

Region 4. Suppose $\lambda_{DM} > x_{\{+,+\}}$. Equilibrium strategies are as follows: The DM never accepts for any h_2 . The advisor's strategy is $\sigma_A = (\varepsilon, \varepsilon, \varepsilon, \varepsilon)$. Given that the DM will never accept for any history of outcomes at the end of period 2, the advisor is indifferent between testing or not.

A.2 Lemma 6 [Hidden Testing Equilibrium]

To characterize the players' expected losses in equilibrium under observable testing, I will use the fact that any equilibrium can be represented as a function ε from the set V of draws by Nature into the set of actions as defined in Section A.1.

Lemma 6 (Hidden Testing on Path)

1. In Region 1, i.e. suppose $x_{\{\emptyset\}} < \lambda_{DM} \leq x_{\{+,-\}}$, then

- if $\lambda_A \leq x_{\{+,-\}}$ the DM accepts iff $v \in \{(+,+), (+,-), (-,+)\}$, and
- if $x_{\{+,-\}} < \lambda_A \leq x_{\{+,+\}}$ the DM accepts iff $v \in \{(+,+)\}$, and
- if $\lambda_A > x_{\{+,+\}}$ the DM never accepts.

2. In Region 2: Define

$$\bar{x} \equiv \frac{q(p_T + p_T(1 - p_T))}{(1 - q)(1 - p_F + p_F(1 - p_F))}.$$

(a) Suppose $\max\{x_{\{\emptyset\}}, x_{\{+,-\}}\} < \lambda_{DM} \leq \bar{x}$, or $x_{\{+,-\}} < \lambda_{DM} \leq x_{\{\emptyset\}}$, then

- if $\lambda_A \leq x_{\{+,-\}}$ the DM accepts iff $v \in \{(+,+), (+,-), (-,+)\}$, and
- if $x_{\{+,-\}} < \lambda_A \leq x_{\{+,+\}}$ the DM accepts iff $v \in \{(+,+)\}$, and
- if $\lambda_A > x_{\{+,+\}}$ the DM never accepts.

(b) Suppose $\bar{x} < \lambda_{DM} \leq x_{\{+\}}$, then

- if $\lambda_A \leq x_{\{+,+\}}$ the DM accepts iff $v \in \{(+,+)\}$, and
- if $\lambda_A > x_{\{+,+\}}$ the DM never accepts.

3. In Region 3, i.e. suppose $x_{\{+\}} < \lambda_{DM} \leq x_{\{+,+\}}$, then

- if $\lambda_A \leq x_{\{+,+\}}$ the DM accepts iff $v \in \{(+,+)\}$, and
- if $\lambda_A > x_{\{+,+\}}$ the DM never accepts.

4. In Region 4, i.e. suppose $\lambda_{DM} > x_{\{+,+\}}$, then the DM never accepts.

Proof of Lemma 6: I will divide the proof into two parts, one for the cases in which the advisor is more enthusiastic than the DM and one for the cases in which he is more reluctant. For each parameter region, I will first propose a profile of strategies σ and a system of beliefs μ , then I will show that this pair (σ, μ) satisfies the definition of equilibrium, and finally, I will show that there is no other equilibrium in which the advisor is strictly better off.

As described in Section 3.1, the advisor prefers accept if and only if her posterior likelihood ratio exceeds λ_A . Hence, his optimal reporting strategy can be written as $\tilde{\sigma}_M : \tilde{H} \rightarrow M$, i.e. the optimal message only depends on the unordered history of outcomes at the end of period 2. In addition, in any equilibrium outlined below, if the DM receives a report of two outcomes then her posterior beliefs satisfy

$$\begin{aligned} \frac{Pr(true|m = \{+, +\})}{Pr(false|m = \{+, +\})} &= x_{\{+, +\}}, \\ \frac{Pr(true|m = \{+, -\})}{Pr(false|m = \{+, -\})} &= x_{\{+, -\}}, \\ \frac{Pr(true|m = \{-, -\})}{Pr(false|m = \{-, -\})} &= x_{\{-, -\}}. \end{aligned}$$

In what follows, I will only specify the DM's beliefs conditional on the advisor disclosing only one outcome or no outcome at all.

Advisor more enthusiastic about accepting than DM. Suppose $\lambda_A < \lambda_{DM}$.

Region 1. Suppose $p_T > p_F$ and $x_{\{\emptyset\}} < \lambda_{DM} \leq x_{\{+, -\}}$.

The following is an equilibrium, referred to as E1: the advisor's testing strategy is

$$\sigma_A = (1 - \varepsilon, 1 - \varepsilon, 1 - \varepsilon, 1 - \varepsilon), \quad (5)$$

the advisor's disclosure strategy is as follows. If $\tilde{h} = \{+, +\}$ send $m = \{+, +\}$ with prob 1; if $\tilde{h} = \{-, -\}$ send $m = \{-, -\}$ with prob ε and $m = \{\emptyset\}$ with prob $1 - \varepsilon$; if $\tilde{h} = \{+, -\}$ send $m = \{+, -\}$ with prob ε and $m = \{+\}$ with prob $1 - \varepsilon$; if $\tilde{h} = \{+\}$ send $m = \{+\}$ with prob 1; if $\tilde{h} = \{-\}$ send $m = \{-\}$ with prob ε and $m = \{\emptyset\}$ with prob $1 - \varepsilon$; if $\tilde{h} = \{\emptyset\}$ send $m = \{\emptyset\}$ with

prob 1. The DM's beliefs for partial disclosure satisfy

$$\frac{Pr(true|m = \{+\})}{Pr(false|m = \{+\})} = \frac{q \left((1-\varepsilon)^3 p_T (1-p_F) + 2\varepsilon (1-\varepsilon) p_T \right)}{(1-q) \left((1-\varepsilon)^3 p_F (1-p_T) + 2\varepsilon (1-\varepsilon) p_F \right)} \quad (6)$$

$$\frac{Pr(true|m = \{-\})}{Pr(false|m = \{-\})} = \frac{q \left((1-\varepsilon)^2 \varepsilon (1-p_T)^2 + 2\varepsilon^2 (1-\varepsilon) (1-p_T) \right)}{(1-q) \left((1-\varepsilon)^2 \varepsilon p_F^2 + 2\varepsilon^2 (1-\varepsilon) p_F \right)} \quad (7)$$

$$\frac{Pr(true|m = \{\emptyset\})}{Pr(false|m = \{\emptyset\})} = \frac{q \left((1-\varepsilon)^3 (1-p_T)^2 + 2\varepsilon (1-\varepsilon)^2 (1-p_T) + \varepsilon^2 \right)}{(1-q) \left((1-\varepsilon)^3 p_F^2 + 2\varepsilon (1-\varepsilon)^2 p_F + \varepsilon^2 \right)} \quad (8)$$

and the DM accepts if and only if $m = \{\{+\}, \{+, -\}, \{+, +\}\}$.

The DM's beliefs are consistent given the advisor's strategy. As $\varepsilon \rightarrow 0$, $\frac{Pr(true|m=\{+\})}{Pr(false|m=\{+\})} \rightarrow x_{\{+, -\}}$, $\frac{Pr(true|m=\{-\})}{Pr(false|m=\{-\})} \rightarrow x_{\{-, -\}}$ and $\frac{Pr(true|m=\{\emptyset\})}{Pr(false|m=\{\emptyset\})} \rightarrow x_{\{-, -\}}$. The DM's strategy is sequentially rational given beliefs since

$$x_{\{-, -\}} < \lambda_{DM} \leq x_{\{+, -\}} < x_{\{+, +\}}. \quad (9)$$

The advisor's disclosure strategy is optimal since for any λ_A such that $x_{\{\emptyset\}} < \lambda_A < \lambda_{DM}$, the DM acts in his interest at any \tilde{h} . For $\lambda_A < x_{\{\emptyset\}}$, the DM does not act in the advisor's interest for $\tilde{h} = \{-, -\}$, but then there is no feasible message for which the DM accepts. Conducting fewer tests never makes the advisor better off.

The equilibrium is an advisor-preferred equilibrium. It maps $v \in \{\{+, -\}, \{-, +\}, \{+, +\}\}$ into *accept*. An equilibrium in which the advisor's expected loss is strictly lower must also map $v = \{-, -\}$ into *accept*. But such an equilibrium cannot exist because for the DM to optimally accept at any v , it must be that the DM accepts at the prior.

Region 2a). Suppose

$$\max \{x_{\{\emptyset\}}, x_{\{+, -\}}\} < \lambda_{DM} \leq \frac{q(p_T + p_T(1-p_T))}{(1-q)(1-p_F + p_F(1-p_F))} \equiv \bar{x},$$

or $x_{\{+, -\}} < \lambda_{DM} \leq x_{\{\emptyset\}}$.

First, for $\lambda_A < x_{\{+, -\}}$, the following is an equilibrium, referred to as E2A: the advisor's testing strategy is given by (5) the advisor's disclosure strategy is as follows. If $\tilde{h} = \{+, +\}$ send $m = \{+, +\}$ with prob ε and $m = \{+\}$ with prob $1 - \varepsilon$; if $\tilde{h} = \{-, -\}$ send $m = \{-, -\}$ with prob ε and $m = \{\emptyset\}$ with prob $1 - \varepsilon$; if $\tilde{h} = \{+, -\}$ send $m = \{+, -\}$ with prob ε and $m = \{+\}$ with prob $1 - \varepsilon$; if $\tilde{h} = \{+\}$ send $m = \{+\}$ with prob 1; if $\tilde{h} = \{-\}$ send $m = \{-\}$ with prob ε and $m = \{\emptyset\}$ with prob $1 - \varepsilon$; if $\tilde{h} = \{\emptyset\}$ send $m = \{\emptyset\}$ with prob 1. The DM's beliefs for partial disclosure

satisfy (7), (8) and

$$\frac{Pr(true|m = \{+\})}{Pr(false|m = \{+\})} = \frac{q \left((1-\varepsilon)^3 (p_T + (1-p_T)p_T) + 2\varepsilon(1-\varepsilon)p_T \right)}{(1-q) \left((1-\varepsilon)^3 (p_F + (1-p_F)p_F) + 2\varepsilon(1-\varepsilon)p_F \right)}. \quad (10)$$

and the DM accepts if and only if $m \in \{\{+\}, \{+, +\}\}$.

The DM's beliefs are consistent given the advisor's strategy. As $\varepsilon \rightarrow 0$, $\frac{Pr(true|m=\{+\})}{Pr(false|m=\{+\})} \rightarrow \bar{x}$, $\frac{Pr(true|m=\{-\})}{Pr(false|m=\{-\})} \rightarrow x_{\{-,-\}}$ and $\frac{Pr(true|m=\{\emptyset\})}{Pr(false|m=\{\emptyset\})} \rightarrow x_{\{-,-\}}$. The DM's strategy is optimal since

$$x_{\{-,-\}} < x_{\{+,-\}} < \lambda_{DM} \leq \bar{x} < x_{\{+,+\}}.$$

If $x_{\{\emptyset\}} < \lambda_A \leq x_{\{+,-\}}$, the advisor's strategy is optimal because the DM acts in his interest at any \tilde{h} . For $\lambda_A < x_{\{\emptyset\}}$, the DM does not act in the advisor's interest for $\tilde{h} = \{-,-\}$, but then there is no feasible message for which the DM accepts. Conducting fewer tests never makes the advisor better off.

The equilibrium is an advisor-preferred equilibrium. It maps $v \in \{\{+, -\}, \{-, +\}, \{+, +\}\}$ into *accept*. An equilibrium in which the advisor's expected loss is strictly lower must also map $v = \{-, -\}$ into *accept*. But such an equilibrium cannot exist because for the DM to optimally accept at any v , it must be that the DM accepts at the prior.

Second, for $x_{\{+,-\}} < \lambda_A < \lambda_{DM} \leq x_{\{\emptyset\}}$ the following is an equilibrium, referred to as E2C: the advisor's testing strategy is given by (5), the advisor's disclosure strategy is to report $m = \tilde{h}$ for all \tilde{h} , the DM believes

$$\frac{Pr(true|m = \{+\})}{Pr(false|m = \{+\})} = x_{\{+\}} \quad (11)$$

$$\frac{Pr(true|m = \{\emptyset\})}{Pr(false|m = \{\emptyset\})} = x_{\{\emptyset\}} \quad (12)$$

$$\frac{Pr(true|m = \{-\})}{Pr(false|m = \{-\})} = x_{\{-\}} \quad (13)$$

and accepts if and only if $m = \{\{+\}, \{+, +\}\}$.

The DM's beliefs are consistent given the advisor's strategy. The DM's strategy is optimal since

$$x_{\{-,-\}} < x_{\{-\}} < x_{\{+,-\}} < \lambda_{DM} \leq x_{\{+\}} < x_{\{+,+\}},$$

the advisor's strategies are optimal since the DM acts in his interest at any \tilde{h} . Conducting fewer tests never makes the advisor better off. The equilibrium is an advisor-preferred equilibrium since it implements the advisor's first-best decision rule, i.e. for any $v \in V$ the equilibrium function e maps into *accept* if and only if the advisor's payoff is maximized by choosing *accept*.

Region 2b). Suppose

$$\bar{x} \equiv \frac{q(p_T + p_T(1 - p_T))}{(1 - q)(1 - p_F + p_F(1 - p_F))} < \lambda_{DM} \leq x_{\{+,+\}}.$$

The following is an equilibrium, referred to as E2B: the advisor's testing strategy is given by (5), the advisor's disclosure strategy is the same as in E2A, the DM's beliefs are the same as in E2A and the DM accepts if and only if $m \in \{+, +\}$.

The DM's beliefs are consistent given the advisor's strategy. As $\varepsilon \rightarrow 0$, $\frac{Pr(true|m=\{+\})}{Pr(false|m=\{+\})} \rightarrow \bar{x}$, $\frac{Pr(true|m=\{-\})}{Pr(false|m=\{-\})} \rightarrow x_{\{-,-\}}$ and $\frac{Pr(true|m=\{\emptyset\})}{Pr(false|m=\{\emptyset\})} \rightarrow x_{\{-,-\}}$. The DM's strategy is optimal since

$$x_{\{-,-\}} < x_{\{+,-\}} < \bar{x} < \lambda_{DM} < x_{\{+,+\}}.$$

The advisor's response is optimal for $x_{\{+,-\}} < \lambda_A < \lambda_{DM}$, the advisor's strategy is optimal because the DM acts in his interest at any \tilde{h} . For $x_{\{-,-\}} < \lambda_A < x_{\{+,-\}}$, the DM does not act in the advisor's interest for $\tilde{h} = \{+, -\}$, but then there is no feasible message for which the DM accepts. Similarly for $\lambda_A < x_{\{-,-\}}$ the DM does not act in the advisor's interest for $\tilde{h} \in \{\{+, -\}, \{-,-\}\}$, but then there is no feasible message for which the DM accepts. Conducting fewer tests never makes the advisor better off.

E2B is an advisor-preferred equilibrium. It cannot be part of an equilibrium that the DM also accepts following $\tilde{h} = \{+, -\}$ or $\tilde{h} = \{-,-\}$. Conditional on $\tilde{h} \in \{\{+, -\}, \{+, +\}\}$ the posterior likelihood ratio is \bar{x} and $\bar{x} < \lambda_{DM}$, and, hence, it is optimal for the DM to reject. Conditional on $\tilde{h} \in \{\{-,-\}, \{+, +\}\}$ the posterior likelihood ratio is even lower than \bar{x} and, hence, it is optimal for the DM to reject.

Region 3. Suppose $x_{\{+,+\}} < \lambda_{DM} \leq x_{\{+,+\}}$. Then E2B is an advisor-preferred equilibrium by the same reasoning as for Region 2b. There cannot be an equilibrium in which the advisor's expected loss is strictly lower, i.e. it cannot be part of an equilibrium that the DM accepts following $\tilde{h} = \{+, -\}$ or $\tilde{h} = \{-,-\}$ since there is no such equilibrium even if λ_{DM} is lower such that the DM's preferences lie in Region 2b.

Region 4. Suppose $\lambda_{DM} > x_{\{+,+\}}$. The following is an equilibrium: the advisor's strategy and the DM's beliefs are as in equilibrium E2C and the DM rejects for any message. I have shown that the DM's beliefs are consistent. As $\varepsilon \rightarrow 0$, the DM's response is optimal since she implements her first-best decision rule, i.e. for any $v \in V$ the equilibrium function e maps into *accept* if and only if the DM's payoff is maximized by choosing *accept*. The advisor's strategy is optimal since the DM's action choice is independent of his strategy. There cannot be an equilibrium in which the advisor is better off since the DM has no reason to deviate from her strategy which yields her first-best decision rule independent of the advisor's strategy.

Advisor more reluctant to accept than DM Suppose $\lambda_A \geq \lambda_{DM}$. The following is an advisor-preferred equilibrium. The advisor's testing strategy is given by (5), his disclosure strategy is as follows. If the advisor prefers to accept, i.e. if $\lambda_A \leq x_{\tilde{h}}$, send $m = \tilde{h}$. Otherwise, i.e. if $\lambda_A > x_{\tilde{h}}$, send $m = \tilde{h}$ with prob ε and $m = \{\emptyset\}$ with prob $1 - \varepsilon$. The DM's beliefs on partial disclosure satisfy

$$\frac{Pr(true|m = \{+\})}{Pr(false|m = \{+\})} = x_{\{+\}},$$

$$\frac{Pr(true|m = \{-\})}{Pr(false|m = \{-\})} = x_{\{-\}},$$

$$Pr(true|m = \{\emptyset\}) = (1 - \varepsilon)^3 \sum_{\tilde{h} \in \{j \in \{+, +, \{+, -\}, \{-, -\} | \lambda_A > x_j\}} Pr(\tilde{h}|true)$$

$$+ 2\varepsilon(1 - \varepsilon)^2 \sum_{\tilde{h} \in \{j \in \{+, \{-\} : \lambda_A > x_j\}} Pr(\tilde{h}|true) + \varepsilon^2,$$

and

$$Pr(false|m = \{\emptyset\}) = (1 - \varepsilon)^3 \sum_{\tilde{h} \in \{j \in \{+, +, \{+, -\}, \{-, -\} | \lambda_A > x_j\}} Pr(\tilde{h}|false)$$

$$+ 2\varepsilon(1 - \varepsilon)^2 \sum_{\tilde{h} \in \{j \in \{+, \{-\} : \lambda_A > x_j\}} Pr(\tilde{h}|false) + \varepsilon^2.$$

The DM rejects $m = \{\emptyset\}$ and for $j \neq \{\emptyset\}$ rejects $m = j$ if and only if $\lambda_{DM} \leq x_j$.

The DM's beliefs are consistent. As $\varepsilon \rightarrow 0$,

$$\frac{Pr(true|m = \{\emptyset\})}{Pr(false|m = \{\emptyset\})} \rightarrow \frac{\sum_{\tilde{h} \in \{j \in \{+, +, \{+, -\}, \{-, -\} | \lambda_A > x_j\}} Pr(\tilde{h}|true)}{\sum_{\tilde{h} \in \{j \in \{+, +, \{+, -\}, \{-, -\} | \lambda_A > x_j\}} Pr(\tilde{h}|false)}.$$

Clearly, the DM's strategy is optimal given $m \neq \{\emptyset\}$. What is left to show is that it is also optimal given $m = \{\emptyset\}$. Her posterior belief about the hypothesis being true based on any on-path message such that $m \neq \{\emptyset\}$ must exceed her prior, i.e.

$$Pr(true) < Pr(true|m \neq \{\emptyset\}),$$

since $m \neq \{\emptyset\}$ if and only if $x_{\{\emptyset\}} < x_{\tilde{h}}$. Since her expected posterior belief is equal to her prior belief, for any on-path message m it must hold that, i.e.

$$Pr(true|m = \{\emptyset\})Pr(m = \{\emptyset\}) + Pr(true|m \neq \{\emptyset\})Pr(m \neq \{\emptyset\}) = Pr(true),$$

this implies that a lack of disclosure must lower her belief about the hypothesis being true:

$$Pr(true) > Pr(true|m = \{\emptyset\}).$$

Since the DM rejects at the prior belief, she must also reject conditional on a lack of disclosure:

$$\frac{Pr(true|m = \{\emptyset\})}{Pr(false|m = \{\emptyset\})} < \frac{Pr(true)}{Pr(false)} < \lambda_{DM}.$$

The advisor's strategy is optimal because he achieves his first-best decision rule, i.e. for any $v \in V$ the equilibrium function e maps into *accept* if and only if the advisor's payoff is maximized by choosing *accept*. This is also the reason why there cannot be another equilibrium in which the advisor is strictly better off.

A.3 Proof of Lemma 1 [Equilibrium]

For this proof I will use the fact that any equilibrium can be represented as a function e from the set V of draws by Nature into the set of actions, as defined in Section A.1. Lemma 5 shows that equilibrium exists and the equilibrium function $e : V \rightarrow \{accept, reject\}$ is unique. Lemma 6 shows that equilibrium exists and that all advisor-preferred equilibria yield the same equilibrium function $e : V \rightarrow \{accept, reject\}$.

A.4 Proof of Theorem 1 and Proposition 1

To compare players' expected losses in equilibrium under hidden and observable testing, I will use the fact that any equilibrium can be represented as a function e from the set V of draws by Nature into the set of actions, as defined in Section A.1. The benchmark case in which player i acquires evidence and takes the decision can also be represented by a function $e : V \rightarrow \{accept, reject\}$, referred to as player i 's *first-best decision rule (henceforth: FB)*. Denote the subset of V for which player i prefers to accept by V_i^* .

Lemma 7 (Hidden vs Observable Testing on Path) *Suppose $\lambda_A > x_{\{+,+\}}$.*

Region 1: *Suppose $x_{\{\emptyset\}} < \lambda_{DM} \leq x_{\{+,-\}}$, then $V_{DM}^* = \{(+,+), (+,-), (-,+)\}$.*

- *If $\lambda_A \leq x_{\{+,-\}}$ then*

$$V_{OT} = V_{HT} = V_{DM}^*,$$

i.e. the DM achieves her FB under both hidden and observable testing, and

- *if $x_{\{+,-\}} < \lambda_A \leq x_{\{+,+\}}$ then*

$$V_{HT} = (+,+) \subset V_{OT} = \{(+,+), (+,-)\} \subset V_{DM}^*,$$

i.e. the DM never achieves her FB, but is closer to FB under observable testing, and

- *if $x_{\{+\}} < \lambda_A \leq x_{\{+,+\}}$ then*

$$V_{HT} = \{\emptyset\} \subset V_{HT} = \{(+, +)\} \subset V_{DM}^*,$$

i.e. the DM never achieves her FB, but is closer to FB under hidden testing, and

Region 2: *Suppose $x_{\{+,-\}} < \lambda_{DM} \leq x_{\{+\}}$, then $V_{DM}^* = \{(+, +)\}$.*

1. *For $x_{\{+,-\}} < \lambda_{DM} \leq \bar{x}$, or $x_{\{+,-\}} < \lambda_{DM} \leq x_{\{\emptyset\}}$, then*

- *if $\lambda_A \leq x_{\{+,-\}}$*

$$V_{DM}^* \subset V_{OT} = \{(+, +), (+, -)\} \subset V_{HT} = \{(+, +), (+, -), (-, +)\},$$

i.e. the DM never achieves her FB, but is closer to FB under observable testing, and

- *if $\lambda_A \leq x_{\{+,-\}}$ then*

$$V_{DM}^* = V_{HT} \subset V_{OT} = \{(+, +), (+, -)\},$$

i.e. the DM achieves her FB under hidden but not under observable testing, and

- *if $x_{\{+,-\}} < \lambda_A \leq x_{\{+,+\}}$ then*

$$V_{DM}^* = V_{HT} = V_{OT},$$

i.e. the DM achieves her FB under both hidden and observable testing, and

2. *If $\bar{x} < \lambda_{DM} \leq x_{\{+\}}$, then*

- *if $\lambda_A \leq x_{\{+,-\}}$ then*

$$V_{DM}^* = V_{HT} \subset V_{OT} = \{(+, +), (+, -)\},$$

i.e. the DM achieves her FB under hidden but not under observable testing, and

- *if $x_{\{+,-\}} < \lambda_A \leq x_{\{+,+\}}$ then*

$$V_{DM}^* = V_{HT} = V_{OT},$$

i.e. the DM achieves her FB under both hidden and observable testing, and

Region 3: *Suppose $x_{\{+\}} < \lambda_{DM} \leq x_{\{+,+\}}$, then $V_{DM}^* = (+, +)$.*

- If $\lambda_A \leq x_{\{+,+\}}$ then

$$V_{DM}^* = V_{HT} = V_{OT},$$

i.e. the DM achieves her FB under both hidden and observable testing, and

Region 4: Suppose $\lambda_{DM} > x_{\{+,+\}}$, then $V_{DM}^* = \{\emptyset\}$. For any λ_A ,

$$V_{DM}^* = V_{HT} = V_{OT},$$

i.e. the DM achieves her FB under both hidden and observable testing.

If $\lambda_A > x_{\{+,+\}}$, then $V_{HT} = V_{OT} = \{\emptyset\}$, i.e. the DM is equally well off under hidden and observable testing and achieves her FB iff $\lambda_{DM} > x_{\{+,+\}}$.

Proof of Lemma 7: Lemma 5 derives the subset of V , denoted by V_{OT} , which is mapped into acceptance under observable testing. Lemma 6 derives the subset of V , denoted by V_{HT} , which is mapped into acceptance under hidden testing. V_i^* contains the draws which lead to a posterior likelihood ratio exceeding preference parameter λ_i (as explained in Section 3.1). \square

If $V_i \subset V_j \subseteq V_{DM}^*$ then the DM must be better off under regime j than regime i , and similarly, if $V_{DM}^* \subseteq V_j \subset V_i$ then the DM must be better off under regime j than regime i , where $i, j \in \{OT, HT\}$ and $i \neq j$. Therefore, the claims about the DM's expected loss in Theorem 1 follow directly from Lemma 7. Proposition 1 follows directly from Lemma 7.

A.5 Corollary 1

By Theorem 1, the DM is weakly better off if and only if both $(\lambda_A, \lambda_{DM}) \notin W_I$ and $(\lambda_A, \lambda_{DM}) \notin W_{II}$. Solve $\max \lambda_{DM} - \lambda_A$ s.t. $(\lambda_A, \lambda_{DM}) \in W_I$. The solution is $\lambda_{DM} - \lambda_A = \bar{x}$. Next, solve $\max \lambda_A - \lambda_{DM}$ s.t. $(\lambda_A, \lambda_{DM}) \in W_{II}$. The solution is $\lambda_A - \lambda_{DM} = x_{\{+\}} - x_{\{\emptyset\}}$. Choose $d = \max \{\bar{x}, x_{\{+\}} - x_{\{\emptyset\}}\}$.

A.6 Proposition 2

The claims follow from Lemmas 5 and 6. If $(\lambda_{DM}, \lambda_A) \in B_I$ then the advisor's FB decision rule is defined by $V_A^* = \{(+, +), (+, -), (-, +)\}$. Hence, he will never achieve his FB, but observable testing yields a higher expected payoff than hidden testing by the same argument as in the proof of Theorem 1. If $(\lambda_{DM}, \lambda_A) \in B_{II}$ then the advisor's FB decision rule is defined by $V_A^* = \{(+, +)\}$. Hence, the advisor achieves his FB under hidden testing but not under observable testing.

A.7 Proof of Lemma 2 [Equilibrium Observable Testing]

Suppose testing is observable, and $x_l < \lambda_{DM} \leq x_{l+1}$ where $l \in \{0, \dots, N-1\}$, and $x_r < \lambda_A \leq x_{r+1}$ where $r \in \mathbb{Z}$. Denote by $x(n)$ the likelihood ratio at the end of period n . Note that the advisor's

expected loss is independent of future test outcomes if these outcomes do not affect the DM's choice.

Part 1.) Choose $\underline{\lambda}_A = x_{2(l+1)-N}$, i.e. $r + 1 = 2(l + 1) - N$. Suppose $\lambda_A > \underline{\lambda}_A$. Then the advisor's optimal strategy is to test with probability $1 - \varepsilon$ in period n if the number of excess positive outcomes is smaller than $l + 1$ at the end of period n , and to test with probability ε in period n if the number of excess positive outcomes is larger or equal to $l + 1$ at the end of period n for $n \in \{0, \dots, N - 1\}$. For any observed h_n , the posterior likelihood ratio equals x_j where j denotes the number of excess positive outcomes in h_n for $n \in \{0, \dots, N\}$. The DM accepts if and only if h_N contains at least $l + 1$ excess positive outcomes.

Beliefs are consistent. Let $\varepsilon \rightarrow 0$. As shown in Section 3, the optimal strategy for the DM is to accept at the end of period N if and only if the likelihood ratio in period N satisfies $\lambda_l \leq x(N)$, i.e. if h_N contains at least $l + 1$ excess positive outcomes.

First, I will show that the advisor strictly prefers to stop testing if the likelihood ratio is equal to x_{l+1} in any period n . This likelihood ratio is only feasible if $n \geq l + 1$. Suppose $x(n) = x_{l+1}$ and $n \geq l + 1$. The lowest possible likelihood ratio in period N is reached if all remaining $N - n$ outcomes are negative and is given by $x(N) = x_{l+1-(N-n)}$. Independent of n , the advisor wants to accept at the lowest possible likelihood ratio since $\lambda_A < \underline{\lambda}_A = x_{2(l+1)-N} \leq x_{l+1-(N-n)}$. Hence, in any period $n \geq l + 1$ where $x(n) = x_{l+1}$ the advisor prefers to accept irrespective of the remaining test outcomes. If he does not test in any future period, then he can guarantee that the DM accepts since $x(N) = x_{l+1} \geq \lambda_{DM}$. If he does test, he cannot guarantee this because it is possible that $x(N) \leq x_l < \lambda_{DM}$ and the DM rejects, e.g. if all remaining outcomes turn out negative.

Second, I will show that the advisor weakly prefers to test in period n if $x(n) \neq x_{l+1}$. Whatever the advisor does at $x(n) > x_{l+1}$ for $n > l + 1$, the likelihood ratio will never fall strictly below x_{l+1} in any future period given that the advisor would stop if the likelihood ratio reached x_{l+1} . Hence, whether the advisor tests or not, the DM's choice is always to accept. If $x(n) < x_{l+1}$, testing affects the DM's choice if and only if eventually $x(N) \geq x_{l+1}$. Since the advisor prefers accept if $x(N) \geq x_{l+1}$ there is an upside but no downside to testing.

Part 2.) Choose $\bar{\lambda}_A = x_k$, where

$$k = \inf \left\{ g \mid g \in \mathbb{Z}_{>0}, g \geq l + 1 + \frac{N - l + 1}{2} \right\}. \quad (14)$$

Suppose $\lambda_A < \bar{\lambda}_A$. Then the advisor's optimal strategy is to test with probability ε in any period n if $x(n) \leq x_l$ and to test with probability $1 - \varepsilon$ in any period n if $x(n) > x_l$ for $n \in \{0, \dots, N - 1\}$. For any observed h_n , the posterior likelihood ratio equals x_j where j denotes the number of excess positive outcomes in h_n for $n \in \{0, \dots, N\}$. The DM accepts if and only if h_N contains at least $l + 1$

excess positive outcomes.

Beliefs are consistent. Let $\varepsilon \rightarrow 0$. The DM's strategy is optimal by the argument of Part 1.). First, I will show that the advisor strictly prefers not to test if $x(n) = x_l$ for any period $n \geq l$. First, note that if at some future period $n' > n$ it holds that $x(n') \geq x_{l+1-(N-n')}$ then in period n' the advisor will be indifferent between testing or not. This is because even if all remaining $N - n'$ outcomes are negative, it holds that $x(N) \geq x_{l+1}$. Hence, the DM will accept regardless of what the advisor does. The highest possible value of the likelihood ratio such that the DM will accept irrespective of future outcomes is reached if all outcomes are positive until $x(n') \geq x_{l+1-(N-n')}$ is satisfied. It is given by $x_{l+1+\frac{N-l}{2}}$ if $N - l$ is even and by $x_{l+1+\frac{N-l+1}{2}}$ if $N - l$ is odd, i.e. the first $l + 1$ outcomes are positive and the majority of the remaining $N - (l + 1)$ outcomes is positive. Since $\bar{\lambda}_A = x_k < \lambda_A$ where k satisfies (14), the advisor would prefer reject even at the highest possible value of the likelihood ratio. Hence, if the advisor were to test at $x(n) = x_l$ in any period $n \geq l$, then provided his testing strategy led the DM to change her decision from reject to accept, the advisor never prefers accept. By contrast, if the advisor did not test at $x(n) = x_l$ for any $n \geq l$ then $x(N) = x_l$ and both the DM and the advisor prefer reject.

Second, if $x(n) > x_l$ the advisor weakly prefers to test. This is because either his testing has no impact on the DM's decision to accept if $x(n) \geq x_{l+1-(N-n)}$ (as argued in Part 1.) or he finds that at some future period $n'' > n$ it holds that $x(n'') = x_l$ and then stops testing and the DM rejects. The advisor benefits if $x(n'') = x_l$ because he wants to reject conditional on the DM rejecting.

Third, if $x(n) < x_l$ the advisor is indifferent between testing or not. Even if he tests for any n such that $x(n) < x_l$ then it must be that $x(N) \leq x_l$ because the advisor stops once the likelihood ratio hits x_l . Hence, the DM rejects regardless of what the advisor does.

In both part 1.) and part 2.) since the DM's strategy is independent of the advisor's strategy and the advisor chooses the optimal action at any given history of outcomes, the equilibrium expected payoff vector is unique.

A.8 Proof of Lemma 3 [Equilibrium Hidden Testing]

Define $l = \{1, \dots, N\}$ to be the exact number of positive outcomes in N tests such that the DM is just willing to accept, i.e.

$$x_{2(l-1)-N} \leq \lambda_{DM} < x_{2l-N}.$$

Define $r = \{1, \dots, N\}$ to be the exact number of positive outcomes in N tests such that the advisor is just willing to accept, i.e.

$$x_{2(r-1)-N} \leq \lambda_A < x_{2r-N}.$$

Part 1.) The following is an equilibrium. The advisor tests with probability $1 - \varepsilon$ in any period n for $n \in \{0, \dots, N-1\}$. His disclosure strategy is as follows. Define a cut-off level \bar{z} by

$$\bar{z} = \inf \left\{ j \mid j \in \mathbb{N}, r \leq j, \lambda_{DM} \leq \frac{q \sum_{s=j}^{s=N} \binom{N}{s} p^s (1-p)^{N-s}}{(1-q) \sum_{s=j}^{s=N} \binom{N}{s} (1-p)^s p^{N-s}} \right\}. \quad (15)$$

If $z \geq \bar{z}$ then he discloses exactly \bar{z} positive and no negative outcomes with probability $1 - \varepsilon$ and he discloses all outcomes with probability ε . If $z < \bar{z}$ he discloses all outcomes with probability 1. If $m = (\bar{z}, 0)$, i.e. if the advisor discloses exactly \bar{z} positive outcomes and no negative outcomes, then the DM's posterior beliefs satisfy

$$x_{m=(\bar{z},0)} = \frac{Pr(\text{true} \mid m = (\bar{z}, 0))}{Pr(\text{false} \mid m = (\bar{z}, 0))}, \quad (16)$$

where

$$\begin{aligned} Pr(\text{true} \mid m = (\bar{z}, 0)) &= q(1-\varepsilon)^{N+1} \sum_{s=\bar{z}}^{s=N} \binom{N}{s} p^s (1-p)^{N-s} \\ &\quad + q \sum_{j=\bar{z}}^{j=N-1} (1-\varepsilon)^{j+1} 1_{j>\bar{z}} \varepsilon^{N-j} \binom{j}{\bar{z}} p^{\bar{z}} (1-p)^{j-\bar{z}}, \end{aligned}$$

$$\begin{aligned} Pr(\text{false} \mid m = (\bar{z}, 0)) &= (1-q)(1-\varepsilon)^{N+1} \sum_{s=\bar{z}}^{s=N} \binom{N}{s} (1-p)^s p^{N-s} \\ &\quad + (1-q) \sum_{j=\bar{z}}^{j=N-1} (1-\varepsilon)^{j+1} 1_{j>\bar{z}} \varepsilon^{N-j+1} \binom{j}{\bar{z}} (1-p)^{\bar{z}} p^{j-\bar{z}}, \end{aligned}$$

and $1_{j>\bar{z}}$ denotes the indicator function, which takes value 1 if $j > \bar{z}$ and 0 otherwise. For any other disclosure $m = (f, g)$, i.e. if the advisor discloses f positive and g negative outcomes, the DM's posterior beliefs satisfy

$$x_{m=(f,g)} = \frac{q \binom{f+g}{f} p^f (1-p)^g}{(1-q) \binom{f+g}{f} p^g (1-p)^{1-f}}. \quad (17)$$

The DM accepts conditional on $m = (\bar{z}, 0)$ and she accepts conditional on $m = (f, g)$ if $f \geq l$. Otherwise, she rejects.

The DM's beliefs are consistent. Suppose $\varepsilon \rightarrow 0$. Then

$$x_{m=(\bar{z},0)} \rightarrow \frac{q \left(\sum_{s=\bar{z}}^{s=N} \binom{N}{s} p^s (1-p)^{N-s} \right)}{(1-q) \left(\sum_{s=\bar{z}}^{s=N} \binom{N}{s} (1-p)^s p^{N-s} \right)}.$$

First, I will show that the DM's strategy is optimal given her beliefs. The DM optimally accepts conditional on $m = (\bar{z}, 0)$ since $\lambda_{DM} \leq x_{m=(\bar{z},0)}$ by the definition of \bar{z} (15). Further, she optimally accepts conditional on $m = (f, g)$ if the number of positive outcomes exceeds l since $\lambda_{DM} < x_{2l-N} \leq x_{2f-N} = x_{m=(f,g)}$. Otherwise, she optimally rejects. Second, the advisor's disclosure strategy is optimal given the DM's strategy. If $z \geq \bar{z}$, he prefers accept since $x_{2z-N} \geq x_{2\bar{z}-N} \geq x_{2r-N} > \lambda_A$ by the definition of \bar{z} (15). Therefore, he optimally reports $m = (\bar{z}, 0)$ if $z \geq \bar{z}$ and the DM accepts. If $z \leq r-1$ he prefers reject. Therefore, he optimally reports $m = (z, g)$ if $z < r-1$ and the DM rejects since $x_{2z-T} \leq x_{2(r-1)-N} \leq \lambda_A < \lambda_{DM}$. If $r \leq z < \bar{z}$, he prefers accept, but for any feasible message $m \in \{(f, g) \mid f \leq z < \bar{z}, g \leq N-z\}$ the DM rejects. It must be weakly optimal for the advisor to conduct N tests because he has the same possible messages available as when he conducted fewer tests and some additional messages.

The above is an advisor-preferred equilibrium. It is optimal for the advisor to pool all outcome realizations with posterior likelihood ratios above some threshold \bar{x} . This is because for both the DM and the advisor the difference between the expected loss from rejecting and the expected loss from accepting increases with the posterior likelihood ratio. Hence, pooling outcomes with posterior likelihood ratios above some threshold \bar{x} makes it most likely that the DM accepts conditional on learning that $x \geq \bar{x}$ and her acceptance in these situations is most valuable to the advisor. The advisor is better off, the lower this threshold \bar{x} subject to the DM accepting conditional on learning $x \geq \bar{x}$ and subject to the advisor preferring acceptance, i.e. $\lambda_A \leq \bar{x}$. The equilibrium outlined above chooses the threshold \bar{x} as low as possible subject to these constraints, where $\bar{x} = x_{2\bar{z}-N}$.

Part 2.) The following is an equilibrium. The advisor tests with probability $1 - \varepsilon$ in any period n for $n \in \{0, \dots, N-1\}$. His disclosure strategy is as follows. Define a cut-off level $\underline{z} \in \mathbb{R}$ by

$$x_{2\underline{z}-N} = \lambda_A. \quad (18)$$

If $z \geq \underline{z}$ then the advisor reports all outcomes with probability 1. Otherwise, the advisor reports $m = (0, 0)$, i.e. no outcomes, with probability $1 - \varepsilon$ and reports all outcomes with probability ε . Conditional on the advisor reporting $m = (0, 0)$, the DM holds the following beliefs:

$$x_{m=0} = \frac{q \left((1 - \varepsilon)^{N+1} \sum_{s=N-\underline{z}+1}^{s=N} \binom{N}{s} p^{N-s} (1-p)^s + \varepsilon^N \right)}{(1-q) \left((1 - \varepsilon)^{N+1} \sum_{s=N-\underline{z}+1}^{s=N} \binom{N}{s} p^s (1-p)^{N-s} + \varepsilon^N \right)}, \quad (19)$$

and for any other report $m = (f, g)$ she holds beliefs (17).

The DM's beliefs are consistent. Suppose $\varepsilon \rightarrow 0$. Then

$$x_{m=0} \rightarrow \frac{q \left(\sum_{s=N-\underline{z}+1}^{s=N} \binom{N}{s} p^{N-s} (1-p)^s \right)}{(1-q) \left(\sum_{s=N-\underline{z}+1}^{s=N} \binom{N}{s} p^s (1-p)^{N-s} \right)}.$$

It is optimal for the DM to accept conditional on $m = (f, g)$ if and only if $f \geq l$ since $x_{m=(f,g)} = x_{2f-N}$ and $x_{2f-N} \geq x_{2(l)-N} > \lambda_{DM}$. In addition, it must be optimal for the DM to reject conditional on $m = (0, 0)$. To see why, partition the space of outcome realizations of the N tests into those for which the number z of positive outcomes satisfies $z \geq \underline{z}$ and those for which $z < \underline{z}$. It must be the case the DM's posterior belief that the hypothesis is true exceeds the prior conditional on $z \geq \underline{z}$:

$$Pr(true) < Pr(true|z \geq \underline{z}).$$

It must also be true that her expected posterior belief is equal to her prior belief, i.e.

$$Pr(true|z < \underline{z}) Pr(z < \underline{z}) + Pr(true|z \geq \underline{z}) Pr(z \geq \underline{z}) = Pr(true).$$

This implies that her posterior belief that the hypothesis is true conditional on $z < \underline{z}$ is lower than the prior:

$$Pr(true) > Pr(true|z < \underline{z}).$$

Since the DM optimally rejects at the prior, she must also optimally reject conditional on $\underline{z} < z$, hence she must reject conditional on $m = (0, 0)$. The advisor's strategy is optimal since his first-best decision rule is implemented. This also implies that there is no equilibrium for which he could achieve a lower expected loss.

A.9 Proof of Proposition 3 and Proposition 5

Part 1. of Proposition 3) Suppose $\lambda_A < x_{N-2}$ and $\lambda_{DM} \in (\bar{x}, x_{N-1}]$ where

$$\bar{x} = \frac{q \sum_{s=N-1}^{s=N} \binom{N}{s} p^s (1-p)^{N-s}}{(1-q) \sum_{s=N-1}^{s=N} \binom{N}{s} (1-p)^s p^{N-s}}, \quad (20)$$

i.e. \bar{x} is the likelihood ratio conditional on $z \geq N-1$ when N tests are run.

Applying Part 1. of the proof of Lemma 2, if $l = N-2$ then $\underline{\lambda}_A = x_{N-2}$. Since $\lambda_A < x_{N-2} = \underline{\lambda}_A$, under observable testing, the advisor stops as soon as the likelihood ratio satisfies $\lambda_{DM} \leq x$, i.e. as soon as $x \geq x_{N-1}$, otherwise he does not stop.

Applying Part 1. of the proof of Lemma 3, with $r = N - 1$, then

$$\bar{z} = \inf \left\{ j \mid j \in \mathbb{N}, N - 1 \leq j, \lambda_{DM} \leq \frac{q \sum_{s=j}^{s=N} \binom{N}{s} p^s (1-p)^{N-s}}{(1-q) \sum_{s=j}^{s=N} \binom{N}{s} (1-p)^s p^{N-s}} \right\} = N,$$

since if $j = N$

$$\frac{q \sum_{s=N}^{s=N} \binom{N}{s} p^s (1-p)^{N-s}}{(1-q) \sum_{s=N}^{s=N} \binom{N}{s} (1-p)^s p^{N-s}} = x_N > \lambda_{DM},$$

but if $j = N - 1$

$$\frac{q \sum_{s=N-1}^{s=N} \binom{N}{s} p^s (1-p)^{N-s}}{(1-q) \sum_{s=N-1}^{s=N} \binom{N}{s} (1-p)^s p^{N-s}} = \bar{x} < \lambda_{DM}.$$

Hence, under hidden testing, the advisor reports all outcomes independent of z and the DM accepts if and only if $z = N$.

The DM's first-best decision rule is to accept if and only if all N outcomes in N tests are positive. If only $N - 1$ outcomes in N tests are positive then she prefers to reject since $x_{N-2} < \lambda_{DM}$. (Note that it is not possible to have $N - 1$ more positive than negative outcomes in N tests.) Under hidden testing, the DM achieves her first-best decision rule. Under observable testing, she does not achieve her first-best decision rule because she accepts if and only if the first $N - 1$ outcomes are positive, irrespective of what the outcome of the last test would be. The advisor must be strictly worse off under hidden testing because he would like to accept if the first $N - 1$ outcomes are positive.

Part 2. of Proposition 3) Suppose $\lambda_A > \lambda_{DM}$ and $\lambda_A > x_k$ where k is given by (14) in the proof of Lemma 2. Under observable testing, Lemma 2 shows that the advisor does not test and the DM always rejects. Under hidden testing, Lemma 3 shows that the advisor runs N tests and the DM accepts if and only if $z \geq \underline{z}$ where \underline{z} solves $x_{2\underline{z}-T} = \lambda_A$. All z for which the DM accepts are z for which she prefers accepting, since if both $\lambda_A > \lambda_{DM}$ and $x_{2z-T} \geq x_{2\underline{z}-T} = \lambda_A$ then also $x_{2z-T} \geq \lambda_{DM}$. Hence, the DM must be strictly better off under hidden testing. The advisor must be strictly better off under hidden testing because his first-best decision rule is implemented.

Proposition 5 follows immediately from the above argument.

A.10 Proof of Proposition 4

Suppose $\lambda_{DM} > \lambda_A$. I will show that the DM achieves her first-best decision rule under hidden testing if $\lambda_{DM} > \bar{x}$, where \bar{x} is given by (20). For $x_N \geq \lambda_{DM} > \bar{x}$, her first-best decision rule is to accept if and only if all N outcomes are positive. To show this, apply Part 1. of the proof of Lemma

3 with $r \leq l$, then

$$\bar{z} = \inf \left\{ j \mid j \in \mathbb{N}, r \leq j, \lambda_{DM} \leq \frac{q \sum_{s=j}^{s=N} \binom{N}{s} p^s (1-p)^{N-s}}{(1-q) \sum_{s=j}^{s=N} \binom{N}{s} (1-p)^s p^{N-s}} \right\} = N.$$

In addition, if $\lambda_{DM} > x_N$ the her first-best decision rule is to reject independent of the N outcomes. There exists an equilibrium in which the advisor tests with probability $1 - \varepsilon$ in every period $n \in \{0, \dots, N-1\}$ and reveals all outcomes with probability 1. Conditional on receiving a report of f positive and g negative outcomes the DM's beliefs are given by (17). The DM rejects independent of the report. Given that the DM achieves her first best decision rule if $\lambda_{DM} > \bar{x}$, she must be weakly better off under hidden testing if $\lambda_{DM} - \lambda_A > \bar{x}$.

Suppose $\lambda_{DM} < \lambda_A$. The proof of proposition 3 shows that the DM is weakly better off under hidden testing if $\lambda_A > x_k$ where k is given by (14). k is increasing in λ_{DM} for $\lambda_{DM} \leq x_N$ and constant for $\lambda_{DM} > x_N$. Define $\tilde{d} \equiv \max_{\lambda_{DM}} x_k(\lambda_{DM}) - \lambda_{DM}$. Hence, the DM is weakly better off under hidden testing if $\lambda_A - \lambda_{DM} > \tilde{d}$. To prove Proposition 4 choose $d = \max \{ \bar{x}, \tilde{d} \}$.

A.11 Proof of Proposition 6 [Delegation]

Part 1.) In the equilibrium under delegation, the advisor implements his first-best decision rule.

Suppose the advisor is more enthusiastic. First, assume the realization of N test outcomes is such that the DM's first-best decision rule maps into acceptance. Then the advisor also prefers to accept since he is more enthusiastic. Therefore, the advisor accepts under delegation. In addition, the DM accepts in equilibrium under hidden testing, because if this were not the case the advisor has a profitable deviation to report all outcomes. By contrast, if the realization of N test outcomes is such that the DM's first-best decision rule maps into rejection, then it may be that the DM rejects in equilibrium under hidden testing, yet the advisor accepts under delegation. This implies that the DM can never be strictly better off under delegation.

Suppose testing is hidden and the advisor is more reluctant. By the proof of Lemma 6 for $N = 2$ and by the proof of Lemma 3 for $N > 2$, a more reluctant advisor also implements his first-best decision rule in equilibrium under hidden testing, which implies that the DM as well off as under delegation.

Part 2.) follows by part 1.) and by Proposition 3.

Part 3.) follows by part 1.) and by Theorem 1.

A.12 Proof of Lemma 4 [DM Commitment]

Suppose the advisor is more enthusiastic. Under observable testing, the DM optimally commits to accept if and only if the advisor has conducted N tests and $\lambda_{DM} \leq x_h^-$. Under hidden testing, the DM

optimally commits to accept if and only if the advisor discloses N tests and $\lambda_{DM} \leq x_{m=\tilde{h}}$. Under observable testing, there is no reason for the advisor to stop testing early because if an additional test outcome is pivotal to the DM's choice then this outcome would lead the DM to accept. If the DM wants to accept the advisor wants to accept. Under hidden testing, there is never a reason to stop testing early, with or without the DM's commitment. The advisor finds it optimal to disclose all outcomes because if hiding an outcome is pivotal to the DM's choice then it must lead the DM to reject when she would have otherwise accepted. This can never be optimal for the advisor by the same reasoning as above.

Suppose the advisor is more reluctant. Under observable testing, the DM optimally commits to reject if and only if the advisor has conducted N tests and $\lambda_{DM} > x_{\tilde{h}}$. Under hidden testing, the DM optimally commits to reject if and only if the advisor discloses N tests and $\lambda_{DM} > x_{m=\tilde{h}}$. An analogous reasoning applies as for the case of a more enthusiastic advisor.

A.13 Proof of Proposition 7 [DM Commitment]

The statement follows by Proposition 4 and by Lemma 7 for $N = 2$ and by Proposition 3 for $N > 2$.

A.14 Proof of Proposition 8 [Advisor Commitment]

The advisor does not strictly benefit from commitment power under observable testing because the DM's optimal strategy is independent of the advisor's optimal strategy and the advisor's equilibrium strategy was constructed using backward induction.

Part 1.) Suppose the advisor is more enthusiastic and has no commitment power. First, assume the realization of the N test outcomes is such that the DM's first-best decision rule maps into acceptance. Then the advisor also prefers to accept since he is more enthusiastic. Therefore, the DM accepts in equilibrium under hidden testing, because if this were not the case then the advisor has a profitable deviation to report more outcomes. Consequently, for this realization of outcomes, the advisor cannot do better with commitment power. By contrast, if the realization of the N test outcomes is such that the DM's first-best decision rule maps into rejection, then it may be that the DM rejects in equilibrium under hidden testing. The advisor is strictly better off with commitment power if there exists some draw of outcomes for which the DM accepts, but for which she would have rejected in the equilibrium under hidden testing without commitment. However, since in the equilibrium under hidden testing without commitment, the DM accepts for all draws for which she prefers to accept, accepting for a larger set of Nature's draws can never make the DM strictly better off.

Part 2.) If the advisor is more reluctant than the DM, then he would indeed commit to the disclosure rule used in the equilibrium under hidden testing without commitment because, as shown in the proof of Lemma 3 for $N > 2$, the advisor achieves his FB in the equilibrium under hidden testing. The rest follows by Proposition 3.

B Online Appendix

B.1 Proof of Proposition 9 [Simultaneous Testing]

Let x_i denote the posterior likelihood ratio conditional on observing i more positive than negative outcomes. Let $x(n)$ denote the posterior likelihood ratio after n tests. Let z denote the total number of positive outcomes.

Part 1.) Suppose testing is observable, and $x_l < \lambda_{DM} \leq x_{l+1}$ where $l \in \{0, \dots, N-1\}$, and $x_r < \lambda_A \leq x_{r+1}$ where $r \in \mathbb{Z}$, and $r+1 \leq l$, i.e. if all tests were observable the advisor would accept for a strictly larger set of outcomes. I will show that if l is even (odd) the advisor chooses the largest odd (even) number of tests n that satisfies $n \leq N$. Hence, given any even (odd) l , if N is odd (even) the advisor chooses $n = N$ and, therefore, the DM can never be strictly better off under hidden testing for any value of r .

As a first step, I will show that if l is even (odd) then the advisor never chooses an even (odd) number of tests. To see why, suppose the advisor has already committed to run n tests, where $n \leq N-1$. An additional test only affects the advisor's expected loss if it affects the DM's choice. There are two situations in which the additional test affects the DM's choice. The first situation is that after n tests the DM rejects, but if the additional test is positive, she accepts. This situation arises if and only if in n tests there were l more positive than negative outcomes, i.e. $x(n) = x_l$. The second situation is that after n tests the DM accepts, but if the additional test is negative, she rejects. This situation arises if and only if in n tests there were $l+1$ more positive than negative outcomes, i.e. $x(n) = x_{l+1}$. Since the advisor prefers accept if there are l more positive than negative outcomes, the advisor is strictly better off with an additional test in the first situation, but strictly worse off in the second situation. If l is even (odd), then the first situation can only arise if n is even (odd) and the second only if n is odd (even). Hence, there is only an upside to an additional test if n is even (odd) and only a downside if n is odd (even). Therefore, the advisor strictly prefers to run an additional test after an even (odd) number of tests.

Consider l is even (odd) and n is odd (even). As a second step, I will show that if the advisor benefits from running an additional two tests at some n then he will benefit from running the largest odd (even) number of tests such that $n \leq N$. There are two situations in which the additional two tests could affect the DM's choice. The first situations is that after n tests the DM rejects, but if both

additional tests are positive, she accepts. This situation arises if and only if in n tests there were $l - 1$ more positive than negative outcomes or a total of $\frac{n+(l-1)}{2}$ positive outcomes, i.e. $x(n) = x_{l-1}$. The second situation is that after n tests the DM accepts, but if both additional tests are negative, she rejects. This situation arises if and only if in n tests there were $l + 1$ more positive than negative outcomes or a total of $\frac{n+l+1}{2}$ positive outcomes, i.e. $x(n) = x_{l+1}$. The advisor will be strictly better off by conducting these test if and only if

$$\lambda_A Pr(false) \left\{ Pr\left(z = \frac{n+l+1}{2} | false\right) p^2 - Pr\left(z = \frac{n+(l-1)}{2} | false\right) (1-p)^2 \right\} + Pr(true) \left[Pr\left(z = \frac{n+(l-1)}{2} | true\right) p^2 - Pr\left(\frac{n+l+1}{2} | true\right) (1-p)^2 \right] > 0 \quad (21)$$

where z denotes the total number of positive outcomes in n tests and

$$Pr(z = j | true) = \binom{n}{j} (1-p)^{n-j} p^j,$$

$$Pr(z = j | false) = \binom{n}{j} (1-p)^j p^{n-j}.$$

Substituting into (21):

$$\begin{aligned} & \lambda_A (1-q) (1-p)^{\frac{n+l+1}{2}} (p)^{\frac{n-(l-1)}{2}} 2 \frac{n!}{\left(\frac{n+l+1}{2} - 1\right)! \left(\frac{n-l-1}{2}\right)!} \left[\frac{1}{(n+l+1)} p - \frac{1}{(n-l+1)} (1-p) \right] \\ & + q (p)^{\frac{n+l+1}{2}} (1-p)^{\frac{n-l+1}{2}} 2 \frac{n!}{\left(\frac{n+l+1}{2} - 1\right)! \left(\frac{n-l+1}{2}\right)!} \left[\frac{1}{n-l+1} p - \frac{1}{n+l+1} (1-p) \right] > 0 \Leftrightarrow \\ & \lambda_A (1-q) (1-p)^l (1-p)^{\frac{n-l+1}{2}} (p)^{\frac{n-l+1}{2}} \left[\frac{1}{(n+l+1)} p - \frac{1}{(n-l+1)} (1-p) \right] \\ & + q p^l p^{\frac{n-l+1}{2}} (1-p)^{\frac{n-l+1}{2}} \left[\frac{1}{n-l+1} p - \frac{1}{n+l+1} (1-p) \right] > 0 \Leftrightarrow \\ & \lambda_A (1-q) (1-p)^l [(n-l+1)p - (n+l+1)(1-p)] \\ & + q p^l [(n+l+1)p - (n-l+1)(1-p)] > 0. \end{aligned} \quad (22)$$

The LHS increases with n since $p > 1-p$. Hence, if (22) holds at some n then it must hold at any larger n .

What is left to show is that the advisor strictly prefers to choose an odd (even) number of tests large enough that it is possible that the outcomes lead the DM to accept. The advisor is indifferent between running $l - 1$ or fewer tests, because even if all $l - 1$ tests were positive, the DM would still reject. However, he must strictly benefit from running $l + 1$ tests, because if the DM accepts than the advisor also prefers to accept. By the argument above, if the advisor benefits from an

additional two tests at $n = l - 1$ then he must benefit from running two additional tests at any larger n . Hence, the advisor chooses the largest odd (even) number of tests feasible.

Part 2.) Suppose testing is observable, and $x_0 < \lambda_{DM} \leq x_1$ and $x_{N-1} < \lambda_A \leq x_N$, i.e. if all tests were observable the DM would accept if and only if the number of positive outcomes strictly exceeds the number of negative outcomes and the advisor would accept if and only if all N outcomes were positive.

I will show that under observable testing, it is optimal for the advisor not to conduct any tests, whereas under hidden testing the advisor will conduct all N tests and disclose outcomes if and only if all N outcomes are positive. Hence, the DM must be strictly better off under hidden testing.

First, suppose testing is observable. I will show that if l is even (odd) then the advisor never chooses an odd (even) number of tests. There are the two situations in which the additional test affects the DM's choice, as described in Part 1.. Since the advisor prefers reject if there is an equal number of positive and negative outcomes, the advisor is strictly better off with an additional test in the situation in which there is one more positive than negative outcome after n tests, i.e. $x(n) = x_1$, but strictly worse off in the situation in which there is an equal number of positive and negative outcomes, i.e. $x(n) = x_0$. Given l even (odd), the first situation can only arise if n is odd (even) and the second only if n is even (odd). Hence, there is only an upside to an additional test if n is odd (even) and only a downside if n is even (odd). Therefore, the advisor strictly prefers to run an additional test after an odd number of tests.

Consider only even n . By an analogous argument as in Part 1.), if the advisor benefits from running an additional two tests at some n then he will benefit from running the largest even number of tests such that $n \leq N$. In particular, there are two situations in which the additional two tests could affect the DM's choice. The first situations is that after n tests the DM rejects, but if both additional tests are positive, she accepts. This situation arises if and only if in n tests there as many positive as negative outcomes, i.e. $x(n) = x_0$. The second situation is that after n tests the DM accepts, but if both additional tests are negative, she rejects. This situation arises if and only if in n tests there were two more positive than negative outcomes, i.e. $x(n) = x_2$. The advisor will be strictly better off by conducting these test if and only if

$$\lambda_A Pr(false) \left\{ Pr\left(z = \frac{n+2}{2} | false\right) p^2 - Pr\left(z = \frac{n}{2} | false\right) (1-p)^2 \right\} + Pr(true) \left[Pr\left(z = \frac{n}{2} | true\right) p^2 - Pr\left(\frac{n+2}{2} | true\right) (1-p)^2 \right] > 0 \quad (23)$$

$$\begin{aligned}
& \lambda_A (1-q) (1-p)^{\frac{n+2}{2}} (p)^{\frac{n}{2}} 2 \frac{n!}{\left(\frac{n+2}{2}-1\right)! \left(\frac{n-2}{2}\right)!} \left[\frac{1}{(n+2)} p - \frac{1}{(n)} (1-p) \right] \\
& + q (p)^{\frac{n+2}{2}} (1-p)^{\frac{n}{2}} 2 \frac{n!}{\left(\frac{n+2}{2}-1\right)! \left(\frac{n}{2}\right)!} \left[\frac{1}{n} p - \frac{1}{n+2} (1-p) \right] > 0 \Leftrightarrow \\
& \lambda_A (1-q) (1-p) [np - (n+2)(1-p)] \\
& + qp [(n+2)p - (n)(1-p)] > 0. \tag{24}
\end{aligned}$$

The LHS increases with n since $p > 1-p$. Hence, if (24) holds at some n then it must hold at any larger n . Therefore, if the advisor runs any tests at all then he will run all N tests.

Given the DM accepts if and only if there are more positive than negative outcomes, the advisor prefers no test to N tests if and only if

$$\begin{aligned}
& Pr(true) < Pr(false) \lambda_A Pr\left(z > \frac{N}{2} | false\right) + Pr(true) Pr\left(z \leq \frac{N}{2} | true\right) \\
& \frac{q Pr\left(z > \frac{N}{2} | true\right)}{(1-q) Pr\left(z > \frac{N}{2} | false\right)} < \lambda_A
\end{aligned}$$

This must hold since for $N > 2$, having exactly $N-1$ positives in N tests is a stronger indication of the hypothesis being true than having strictly more positive than negative outcomes:

$$\frac{q Pr\left(z > \frac{N}{2} | true\right)}{(1-q) Pr\left(z > \frac{N}{2} | false\right)} < x_{N-1} \equiv \frac{qp^{N-1}}{(1-q)(1-p)^{N-1}},$$

and $x_{N-1} < \lambda_A$.

Finally, suppose testing is hidden. Then the advisor would commit to run all N tests, because even if testing is sequential and his strategy space is richer, there is no downside to choosing to run N tests. By the same argument as in Lemma 3, the advisor fully discloses outcomes if and only if he himself would accept, i.e. he fully discloses if and only if $x(N) = x_N$ since $x_{N-1} < \lambda_A \leq x_N$.

B.2 Infinite Horizon with Testing Cost

The following modifications are made to the model in Section 2: There are infinitely many discrete periods. The DM's payoff function is unchanged, but the advisor's payoff function is modified to include a cost $c > 0$ for each test run. Under hidden testing, in each period n , the advisor first privately chooses to test ($a = 1$) or not ($a = 0$). In addition, at the end of period n , he chooses to send a message ($r = m$) or not ($r = 0$), where $m \in M_n$ and the message space is the unordered history at the end of period n , i.e. $M_n = \mathcal{P}(\tilde{h}_n)$. After receiving a message, the DM chooses $\tau \in \{accept, reject\}$. The DM does not know in which period the message was sent.²⁴ A testing

²⁴This assumption is made to insure that the DM does not know how many tests have been run.

strategy for the advisor is $\sigma_A : H_n \rightarrow \{0, 1\}$ for $n = 0, 1, \dots$. A reporting strategy for the advisor is $\sigma_M : H_n \rightarrow 0 \times M_n$ for $n = 1, 2, \dots$. A strategy for the DM is $\sigma_A : M \rightarrow \{accept, reject\}$, where $M = \cup_{n=1}^{\infty} M_n$. The solution concept is an advisor-preferred sequential equilibrium. Under observable testing, in each period n , the advisor chooses publicly whether to test or not and then the DM chooses whether to wait ($\tau = 0$), or to accept ($\tau = accept$) or to reject ($\tau = reject$). A strategy for the advisor is $\sigma_A : H_n \rightarrow \{0, 1\}$ for $n = 0, 1, \dots$. A strategy for the DM is $\sigma_A : H_n \rightarrow \{0, accept, reject\}$. The solution concept is a sequential equilibrium.²⁵ Any unordered history \tilde{h}_n can be characterized by $\tilde{h}_n = (x, y)$, where x denotes the number of positive outcomes and y denotes the number of negative outcomes at the end of period n . Denote number of excess positive outcomes by $k = x - y$ where $k \in \mathbb{Z}$. Denote any unordered history with k excess positive outcomes by h^k and denote the posterior likelihood ratio by $\frac{Pr(true|h^k)}{Pr(false|h^k)} \equiv x_k$.

Insurance Effect First, I will argue that the insurance effect does not exist. Under observable testing, the DM's optimal strategy is to wait if and only if the advisor tested in the current period. Otherwise, the DM optimally rejects if and only if there are l or fewer excess positive outcomes, where $l = \sup \{k \in \mathbb{Z} | x_k < \lambda_{DM}\}$, independent of the current period. Since the number of excess positive outcomes is sufficient to determine the DM's final decision and to form posterior beliefs about the state, the advisor's optimal testing strategy is a function only of the number of excess positive outcomes. Under hidden testing, the advisor's first-best decision rule will be implemented in equilibrium by the same reasoning as used in the proof of Lemma 3. This implies that if the advisor stops testing when there are r or fewer excess positive outcomes then the DM will reject and otherwise she will accept, where $r = \sup \{k \in \mathbb{Z} | x_k < \lambda_A\}$. Hence, the advisor's optimal testing strategy is again a function only of the number of excess positive outcomes.

In order for hidden testing to be beneficial for the DM, there must be some realizations of Nature's draws such that acceptance is chosen in equilibrium. A necessary condition for acceptance to be chosen in equilibrium is that the advisor tests at any history where the number of excess positive outcomes k satisfies $0 \leq k \leq r$, i.e. for any k satisfying $0 \leq k \leq r$ the advisor's continuation value of testing exceeds his continuation value of stopping. Suppose this were the case. Under observable testing, his continuation value of stopping at k excess positive outcomes is the same as under hidden testing for $k \leq l$ or $k > r$, and it is lower for $l < k \leq r$. Therefore, the advisor must also prefer testing to not testing at k excess positive outcomes satisfying $0 \leq k \leq r$ when testing is observable. Hence, the insurance effect cannot exist.

²⁵In this model, multiple equilibria exist under observable testing. To select a unique equilibrium I assume that the advisor tests in the current period if he is indifferent between testing in the current or in a later period and that the DM takes her final decision in the current period if she is indifferent between taking the final decision in the current or in a later period.

Skepticism Effect Define two thresholds of c :

$$\bar{c} = \min \left\{ \frac{q(1-p)p}{q(1-p) + (1-q)p}, \frac{3p^2q - 3p^3q + p^4q}{2 - p^3 - q + 3pq - 3p^2q + 2p^3q} \right\},$$

$$\underline{c} = \max \left\{ \frac{q(1-p)p^2}{2q(1-p) + 2(1-q)p}, \frac{q(1-p)^2p}{q(1-p)^2 + (1-q)p^2} \right\}.$$

I will show that for any q , there exists a p such that $\bar{c} > c > \underline{c}$. In addition, if cost c satisfies $\bar{c} > c > \underline{c}$ then I will show that there is a combination of preference parameters λ_A and λ_{DM} such that the skepticism effect exists.

At $p = 1$, $\bar{c} = \underline{c} = 0$. As $p \rightarrow 1$, $\underline{c} = \frac{q(1-p)p^2}{2q(1-p) + 2(1-q)p}$ and $\bar{c} = \frac{q(1-p)p}{q(1-p) + (1-q)p}$, where $\bar{c} > \underline{c}$ since both $\frac{\partial \bar{c}}{\partial p} < 0$ and $\frac{\partial \underline{c}}{\partial p} < 0$ but $-\frac{\partial \bar{c}}{\partial p} > -\frac{\partial \underline{c}}{\partial p}$. Hence, $\bar{c} > \underline{c}$ for p sufficiently large.

Suppose p is sufficiently large, the cost satisfies $\bar{c} > c > \underline{c}$, and $\lambda_A = 0$ and

$$\bar{x} \equiv \frac{q(p + p(1-p))}{(1-q)((1-p) + (1-p)p)} < \lambda_{DM} < \frac{q(p^2 + p^2(1-p) + p^2(1-p)^2)}{(1-q)((1-p)^2 + (1-p)^2p + p^2(1-p)^2)} \equiv \hat{x}.$$

Under observable testing, the DM's optimal strategy is to wait if and only if the advisor tested in the current period. Otherwise, the DM optimally accepts if and only if at least one excess positive outcome was found, independent of the current period. Since the number of excess positive outcomes is sufficient to determine the DM's final decision and to form posterior beliefs about the state, the advisor's optimal testing strategy is a function only of the number of excess positive outcomes. The advisor optimally stops testing when he has found one or more excess positive outcomes, since the DM accepts and this is the advisor's preferred action independent of the state. In addition, the advisor optimally stops testing when he has found one or more excess negative outcomes. Given his posterior beliefs after finding one excess negative outcome, the expected cost of testing outweighs the expected benefit. In particular, even if all future test outcomes were favorable for the advisor then two tests are needed to convince the DM to accept, but the advisor is not willing to pay for two tests since

$$\begin{aligned} -Pr(\text{true}|h^{-1})(1-p^2) - 2c &< -Pr(\text{true}|h^{-1}) \\ Pr(\text{true}|h^{-1})p^2 &< 2c \\ \frac{q(1-p)p^2}{q(1-p) + (1-q)p} &< 2c, \end{aligned} \tag{25}$$

which is implied by $c > \underline{c}$. Finally, given that the advisor optimally stops after either one excess positive or one excess negative outcome, he optimally tests when he has as many positive as

negative outcomes since

$$\begin{aligned} -Pr(\text{true}|h^0)(1-p) - c &> -Pr(\text{true}|h^0) \\ Pr(\text{true}|h^0)p &> c \\ qp &> c, \end{aligned}$$

which is also implied by $c > \underline{c}$. Next, suppose testing is hidden. I will show that it cannot be part of an equilibrium for the DM to accept if and only if at least one positive outcome was reported. To see why, suppose the DM were to accept if and only if at least one positive outcome was reported. The advisor's optimal reporting strategy is to hide any negative outcomes and to report any positive outcomes whenever he has stopped testing. The advisor optimally stops testing when he has found at least one positive outcome (independent of how many negative outcomes he has found), because then the DM acts in the advisor's interest. In addition, the advisor optimally stops testing when he has found two negative outcomes, because his expected cost of testing outweighs his expected benefit. In particular, even if all future test outcomes were favorable for the advisor then one test is needed to convince the DM, but the advisor is not willing to pay for even one test since

$$\begin{aligned} -Pr(\text{true}|h^{-2})(1-p) - c &< -Pr(\text{true}|h^{-2}) \\ Pr(\text{true}|h^{-2})p &< c \\ \frac{q(1-p)^2 p}{q(1-p)^2 + (1-q)p^2} &< c, \end{aligned} \tag{26}$$

which is also implied by $c > \underline{c}$. By contrast, the advisor optimally continues testing when he has found one negative outcome since

$$\begin{aligned} -Pr(\text{true}|h^{-1})(1-p) - c &> -Pr(\text{true}|h^{-1}) \\ Pr(\text{true}|h^{-1})p &> c \\ \frac{q(1-p)p}{q(1-p) + (1-q)p} &> c, \end{aligned} \tag{27}$$

which is implied by $\bar{c} > c$. If he continues testing after one negative, then he also optimally starts testing since he only requires one additional positive outcome to convince the DM and his posterior belief that the state is true is higher than in (27). The advisor's optimal testing and reporting strategies imply that the DM observes a report of one positive outcome either if the first outcome was positive, or if the first outcome was negative and the second positive. The DM optimally rejects such a report given $\bar{x} < \lambda_{DM}$. Next, I will show that there exists an equilibrium in which the DM accepts if and only if at least two positive outcomes were reported. Suppose the DM followed

this strategy. The advisor then optimally hides any negative outcomes and reports any positive outcomes whenever he stops testing. In addition, the advisor optimally stops testing when he has found at least two positive outcome, because then the DM acts in the advisor's interest. Further, the advisor optimally stops testing if he has found one negative and no positive outcomes. This is because even if all future test outcomes were favorable for the advisor then two tests are needed to convince the DM, but the advisor is not willing to pay for two tests since (25) holds. In addition, he optimally stops if he has found three negatives and one positive outcome. This is because even if all future test outcomes were favorable for the advisor then one test is needed to convince the DM, but the advisor is not willing to pay for one test since (26) holds. However, it is optimally for him to start testing since

$$\begin{aligned}
& -q \left(1 - p + p(1 - p)^3 \right) \\
& - (q(1 - p) + (1 - q)(p)) c \\
& - \left(qp^2 + (1 - q)(1 - p)^2 \right) 2c \\
& - \left(q(1 - p)p^2 + (1 - q)(1 - p)^2 p \right) 3c \\
& - \left(q(1 - p)^2 p + (1 - q)(1 - p)p^2 \right) 4c > -q \\
& \Leftrightarrow \\
& \frac{3p^2q - 3p^3q + p^4q}{2 - p^3 - q + 3pq - 3p^2q + 2p^3q} > c \tag{28}
\end{aligned}$$

which is implied by $\bar{c} > c$. Given (28), it must also be optimal to continue testing if he has one positive and one negative or if he has one positive and no negative. In either case, he only needs one additional positive outcome and his posterior belief of the state being true is at least as high as when he started testing. Finally, he will continue testing after two negative outcomes and one positive outcome since (27) holds. This implies that the DM accepts conditional on $h_2 = (+, +)$ or $h_3 = (+, -, +)$ or $h_4 = (+, -, -, +)$. In equilibrium, the DM's posterior likelihood ratio conditional on observing two positive outcomes is given by \hat{x} and since $\lambda_{DM} < \hat{x}$, it is indeed optimal for the DM to accept. Hence, in the equilibrium under hidden testing the DM accepts at a higher likelihood ratio than in the equilibrium under observable testing since

$$x_1 \equiv \frac{qp}{(1 - q)(1 - p)} < \hat{x}.$$

References

- P. Aghion and J. Tirole. Formal and real authority in organizations. *Journal of Political Economy*, 105(1):1–29, 1997.
- R. Argenziano, S. Severinov, and F. Squintani. Strategic information acquisition and transmission. *American Economic Journal: Microeconomics*, 8(3):119–55, August 2016.
- I. Brocas and J. Carrillo. Influence through Ignorance. *The RAND Journal of Economics*, 38(4): 931–947, 2007.
- Y.-K. Che and N. Kartik. Opinions as Incentives. *Journal of Political Economy*, 117(5):815–860, 2009.
- A. J. Davis. Prosecutors Who Intentionally Break the Law . *Criminal Law Brief*, 1(1):16, 2006.
- M. DeGroot. New York: McGraw-Hill, 1970.
- I. Deimen and D. Szalay. Information, Authority, and Smooth Communication in Organizations. CEPR Discussion Papers 10969, C.E.P.R. Discussion Papers, Dec 2015.
- W. Dessein. Authority and communication in organizations. *The Review of Economic Studies*, 69 (4):811–838, 2002.
- A. Di Tillio, M. Ottaviani, and P. N. Sorensen. Strategic Sample Selection. CEPR Discussion Papers 12202, C.E.P.R. Discussion Papers, Aug 2017.
- R. Dur and O. Swank. Producing and Manipulating Information. *The Economic Journal*, 115 (500):185–199, 2005.
- S. Epstein. *Impure Science: AIDS, Activism, and the Politics of Knowledge*. University of California Press, 1998.
- J. Farrell and J. Sobel. Voluntary Revelation of Product Information. *Unfinished Mimeo*, 1983.
- M. Felgenhauer and P. Loerke. Bayesian Persuasion with Private Experimentation. *International Economic Review*, 58:829–856, 2017.
- M. Felgenhauer and E. Schulte. Strategic Private Experimentation. *American Economic Journal: Microeconomics*, 6(4):74–105, 2014.
- D. Gerardi and L. Yariv. Costly Expertise. *American Economic Review*, 98(2):187–93, May 2008.

- B. Goldacre, S. Lane, K. Mahtani, C. Heneghan, I. Onakpoya, I. Bushfield, and L. Smeeth. Pharmaceutical companies' policies on access to trial data, results, and methods: audit study. *BMJ*, 358, 2017.
- S. Grossman. The Informational Role of Warranties and Private Disclosure about Product Quality. *The Journal of Law & Economics*, 24(3):461–483, 1981.
- E. Henry. Strategic Disclosure of Research Results: The Cost of Proving Your Honesty. *The Economic Journal*, 119(539):1036–1064, 2009.
- E. Henry and M. Ottaviani. Research and the Approval Process: The Organization of Persuasion. CEPR Discussion Papers 11939, C.E.P.R. Discussion Papers, Mar 2017.
- E. Kamenica and M. Gentzkow. Bayesian Persuasion. *American Economic Review*, 101(6):2590–2615, October 2011.
- H. Li and W. Suen. Delegating Decisions to Experts. *Journal of Political Economy*, 112(S1): S311–S335, 2004.
- J. Liebman, J. Fagan, and V. West. "a broken system: Error rates in capital cases, 1973-1995". *Columbia Law School, Public Law Research Paper*, (15), June 2000.
- S. Matthews and A. Postlewaite. Quality Testing and Disclosure. *The RAND Journal of Economics*, 16(3):328–340, 1985.
- P. Milgrom. Good News and Bad News: Representation Theorems and Applications. *The Bell Journal of Economics*, 12(2):380–391, 1981.
- P. Milgrom and J. Roberts. Price and Advertising Signals of Product Quality. *Journal of Political Economy*, 94(4):796–821, 1986.
- A. Pollack. 'Viagra for Women' is Backed by an F.D.A. Panel. *The New York Times*, page A1, June 2015.
- T.C. Schelling. *The Strategy of Conflict*. Harvard University Press, 1960.
- S. Shavell. Acquisition and Disclosure of Information Prior to Sale. *The RAND Journal of Economics*, 25(1):20–36, 1994.
- H. Shin. The Burden of Proof in a Game of Persuasion. *Journal of Economic Theory*, 64(1): 253–264, 1994.