Discussion Paper Series – CRC TR 224

Discussion Paper No. 145
Project B 04

Hidden Testing and Selective Disclosure of Evidence

Claudia Herresthal*

January 2020

*University of Bonn, Adenauerallee 24-42, 53113 Bonn, Germany (email: claudia.herresthal@unibonn.de)

Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 is gratefully acknowledged.

Hidden Testing and Selective Disclosure of Evidence*

Claudia Herresthal†

January 16, 2020

Abstract

An agent can sequentially run informative tests about an unknown state and disclose (some or all) outcomes to a decision maker who then faces an approval choice. Players agree on the optimal choice under certainty, but the decision maker has a higher approval threshold than the agent. I compare the case where testing is hidden and the agent chooses which test outcomes to verifiably disclose to the case where testing is observable. When testing is observable, I show that the agent may strategically stop testing even if further tests could yield a mutual benefit. I find conditions under which the decision maker is strictly better off under hidden testing and in some equilibria both players are strictly better off under hidden testing than in the unique equilibrium under observable testing.

Keywords: endogenous information acquisition, verifiable disclosure, transparency, questionable research practices

JEL codes: D83, D82

*The latest version of this paper can be found at https://cherresthal.github.io/CHerresthal_HiddenTesting.pdf. I would like to thank Leonie Baumann, Inga Deimen, Francesc Dilme, Matt Elliott, Robert Evans, Pawel Gola, Ole Jann, Helene Mass, Meg Meyer, Marco Ottaviani, Anna Sanktjohanser, Dezso Szalay, Adrien Vigier and conference and seminar audiences at Conference on Economic Design, Game Theory Conference at Stony Brook, EEA, TSE, Bocconi, BI Norwegian Business School, University of Arizona, Birkbeck, Münster, Cambridge and Oxford for valuable comments and feedback. Funding by the German Research Foundation (DFG) through CRC TR 224 (Project B04) is gratefully acknowledged.

†University of Bonn, Adenauerallee 24-42, 53113 Bonn, Germany (email: claudia.herresthal@uni-bonn.de).

Pharmaceutical companies have recently come under scrutiny for selective reporting of clinical trial outcomes.[1] As a response to demands for greater transparency, several companies have pledged to register trials and report their outcomes in open online databases (Goldacre et al. [2017]). At first sight, it seems that such transparency would improve regulation because pharmaceutical companies can no longer hide trials with unfavorable outcomes. However, companies may also strategically respond by running fewer trials and this could mean that the regulator has to base his approval decision on weaker evidence.

This paper analyzes how transparent information acquisition affects the likelihood of correct approval decisions and welfare. A decision-maker (regulator) has to take an approval decision based on evidence gathered by an agent (company). Players differ in how much they value taking an appropriate action in each state, e.g. the regulator may be relatively more averse to approving an unsafe product than the company. Consequently, players have different threshold beliefs at which they favor approval, e.g. the regulator needs to hold a higher belief that the product is safe than the company in order to favor approval.

I first consider a setting of *observable testing*, which corresponds to having a trial registry. The agent can sequentially choose to run tests over a finite number of periods. The outcomes of these tests are publicly observed and each outcome is either evidence in favor of the product being safe or unsafe. In the final period, the decision maker (DM) chooses to approve or not. I contrast this with a setting of *hidden testing*, which corresponds to not having a trial registry. In this case, the DM neither observes how much the agent has tested nor what he has found. In the final period, the agent strategically chooses which outcomes to disclose to the DM, that is, the agent can hide but not forge outcomes. The DM takes an approval decision based on the disclosed outcomes.

One might expect that the DM is better off when testing is observable than when testing is hidden. After all, when testing is observable, the DM has the advantage that he can

[1]For example, GlaxoSmithKline was accused of having withheld data that suggested that its antidepressant Paxil was linked to suicidal behavior in teenagers (Rabin [2017]), and the Cochrane Review concluded that Roche withheld trial data on its influenza drug Tamiflu to make the drug look more effective (Jefferson et al. [2014]).

react optimally to the evidence gathered by the agent. However, I find parameter conditions under which the DM is strictly better off in any equilibrium under hidden testing than in the unique equilibrium under observable testing. What is more, under these conditions some equilibria under hidden testing yield a strictly higher payoff to both players than the unique equilibrium under observable testing.

When testing is observable, the DM can react optimally to the evidence gathered, but the agent will acquire evidence strategically. In particular, the agent can manipulate the DM's choice by stopping to test following certain outcomes, while continuing to test following other outcomes. I characterize the equilibrium profile of strategies (Proposition 1) and show that the agent may stop testing even if stopping results in approval and further tests could lead *both* players to prefer rejection (Lemma 1). The reason is that, by running further tests, the agent risks that the additional outcomes convince the DM to reject, while the agent continues to prefer approval. Therefore, having a mutual interest in learning is not sufficient for the agent to continue testing.

When testing is hidden, the agent has no downside from gathering evidence, but he is strategic in which outcomes he reveals to the DM. In equilibrium, the DM may approve conditional on what the agent reveals, even though the DM would have optimally rejected had he been able to observe all of the evidence gathered (Proposition 2).

I show that both players can be strictly better off under hidden testing than under observable testing. I first characterize conditions under which rejection is never chosen in the unique equilibrium under observable testing, but rejection is chosen in any equilibrium under hidden testing for some outcome realization (Theorem 1). When these conditions hold, the DM is strictly better off under hidden testing. What is more, in the agent-preferred equilibrium under hidden testing both players are strictly better off than in the unique equilibrium under observable testing (Corollary 2). In addition, if the agent had the power to commit to a disclosure strategy ex ante, both players would be better off under hidden compared to observable testing (Lemma 3).

The conditions in Theorem 1 specify that players have a mutual interest in learning, yet the conflict is large enough such that under observable testing the agent stops testing before he can find any evidence that would lead the DM to reject. When testing is observable, there is an upside to testing for the agent due to the mutual interest in learning. However, the agent stops because if he were to continue testing, it is too likely that the DM will reject, yet the agent will not be convinced himself that rejection is the optimal choice. By contrast, when testing is hidden, mutual gains from testing can be realized. In particular, the agent tests and reveals at most some threshold number of outcomes in favor of approval. This threshold number is such that the DM is just willing to approve conditional on inferring that at least as many outcomes as revealed were in favor of approval. Therefore, there are some outcome realizations following which the DM approves whereas he would have rejected under observable testing, and this is what motivates the agent to test when testing is hidden. The DM is better off under hidden testing because he at least learns if the evidence lead the agent to prefer rejection.

In the context of the leading example, suppose the product is believed to be safe with a sufficiently high likelihood that the regulator would approve without further tests. Without a trial registry, the company may privately run additional tests and withdraw their application for approval if and only if it finds strong evidence that the product is unsafe. However, with a trial registry, the company is discouraged from performing such tests. The reason is that further tests could yield weak evidence that the product is unsafe. Then the company would still want to seek approval, but the regulator would reject.

The key difference to existing work on the comparison between hidden and observable testing is the assumption that the two players agree on the optimal action at extreme beliefs but disagree at intermediate beliefs. This is crucial for the novel insight into how both players can benefit from hidden testing. In closely related work, Henry [2009] and Felgenhauer and Loerke [2017] also find that the DM is better off in any Pareto-undominated equilibrium under hidden testing but, unlike this paper, they find that the agent is worse off. In Section

6, I find parameter configurations for my setting under which only the DM benefits from hidden testing for reasons related to those identified by their work. Then I explain how these reasons differ from the ones which lead to the novel insight that both players can benefit from hidden testing.

1 Related Literature

My work builds on the extensive literature on strategic information acquisition and verifiable disclosure.[2] What differentiates my set-up from those in the related literature is that the players have different approval thresholds, which implies that the difference between the two players' optimal actions does not vary monotonically with the belief.[3] This is a necessary feature for the main result (Theorem 1 and Corollary 2) about how hidden testing can help players to realize mutual gains from information acquisition.

Some of the existing work on the comparison between hidden and observable testing assumes that the agent can acquire at most one signal. Matthews and Postlewaite [1985] study a seller who can acquire a single costless signal about product quality.[4] When disclosure is compulsory, the seller does not test and buyers take his claimed ignorance at face value. But when disclosure is voluntary, the seller tests and reveals the signal if the quality is good, which allows buyers to learn about quality. Dahm et al. [2009] study disclosure rules, assuming a company can decide whether or not to run a single trial whose outcome is either positive, negative or inconclusive. They find that compulsory registries combined with a voluntary results database can result in full transparency, but reduce the company's incentive to test. Gall and Maniadis [2019] analyze a tournament between researchers with and without certain

[2] A large literature studies persuasion when the agent is exogenously informed, starting with Milgrom [1981], Milgrom and Roberts [1986].

[3] In the related literature, either the agent prefers acceptance irrespective of the state so that players can only ever agree on acceptance (Matthews and Postlewaite [1985], Dahm et al. [2009], Felgenhauer and Loerke [2017], Brocas and Carrillo [2007], Felgenhauer and Schulte [2014], Di Tillio et al. [2017a], Henry and Ottaviani [2019], Janssen [2018]), or the agent's ideal action always differs from the DM's ideal action by a constant independent of the state (Henry [2009], Che and Kartik [2009]).

[4] See also Farrell and Sobel [1983] and Shavell [1994].

possibilities of inflating outcomes at a cost and find that preventing researchers from selective reporting also discourages more severe questionable research practices such as fabrication. Di Tillio et al. [2017b] study how manipulation of the design and reporting of medical trials by a biased researcher can influence and, in some circumstances, improve informativeness.

In my set-up, the agent acquires information by running a given test sequentially, which is also assumed by Felgenhauer and Schulte [2014] for studying persuasion under hidden testing and by Brocas and Carrillo [2007], and in continuous time by Henry and Ottaviani [2019] and McClellan [2017], for studying persuasion under observable testing.

The most closely related papers on the comparison between hidden and observable testing share the assumption that an agent can acquire multiple signals from a given test, which expands the agent's scope for strategically influencing the DM's choice. My paper shares with Henry [2009] and Felgenhauer and Loerke [2017] the conclusion that the DM benefits from hidden testing. However, I show that both players can benefit from hidden testing, whereas they show that the agent is worse off under hidden testing.[5]

Henry [2009] assumes that the agent commits ex ante to a quantity of costly research, which maps into a state-dependent number of infinitesimal positive and negative signals. If his choice is hidden, the agent conducts more research to have a larger pool of signals to select from but interestingly, in the end, this leaves the DM better informed since he can infer all signals due to unraveling in the vein of Milgrom [1981] and Grossman [1981].[6] Felgenhauer and Loerke [2017] study an agent who tests sequentially and can decide how informative each test will be. Surprisingly, they find that the agent runs only a single test in any Pareto-undominated equilibrium, whether testing is observable or hidden. However, if testing is hidden the agent runs a more informative test, because this makes it credible that he will not run further tests even if the outcome is unfavorable.

Janssen [2018] also compares hidden and observable testing in the context of a signaling

[5] In Section 6, I give parameter conditions for which the DM benefits and the agent suffers from hidden testing in my setting and discuss the differences to existing work in more detail.

[6] In my setting, the game of hidden testing has no unique sequential equilibrium and the DM does not necessarily infer all the agent's findings in equilibrium.

game. An agent who is perfectly informed about the state draws test outcomes to convince the DM to approve. When testing is hidden, the DM draws inferences about the agent's prior knowledge based on his decision to start testing or not. The DM always weakly prefers observable testing, while the agent prefers observable testing when the DM is close to accepting without evidence.

Di Tillio et al. [2017a] compare the DM's payoff under two scenarios, one in which she allows a biased agent to collect a sample of a given size in private and report a fixed number of draws and one in which she restricts him to collect the same fixed number of draws in public. Their focus lies on identifying properties of the distribution of signals which imply that more selective disclosure may leave the DM either better or worse off. By contrast, my focus lies on identifying how the conflict of interest between players affects the DM's benefit from hidden testing when the agent can decide the number of disclosed signals strategically.

Che and Kartik [2009] do not compare different informational environments, but study the ideal bias of an agent who in private exerts costly effort to increase his chances of observing a single normally distributed signal.[7] They show that the DM prefers a biased to an unbiased agent provided his bias is sufficiently large. A more biased agent does not report his signal for a larger range of realizations but has greater incentives to exert effort. In my setting, I show that making testing hidden instead of observable can motivate the agent to acquire more signals and, although the agent will be able to suppress some of those signals, the DM is made better off. The increased motivation to test stems from the possibility to exploit a mutual gain from information acquisition, not from the need to bridge a larger conflict of interest.

2 Model

Two players, a decision-maker (DM) and an agent (A), are uninformed about the state of the world $\omega \in \{\omega_L, \omega_R\}$ and share a prior belief $q_0 = Pr(\omega_R) \in [0, 1]$. Time is discrete and

[7] See also Gerardi and Yariv [2008] and Dur and Swank [2005].

finite, $t = 1, \ldots, T$.[8] At the end of period T, the DM chooses an action $a \in \{a_L, a_R\}$ which affects the payoff of player i as follows,[9]

$u_i(a, \omega)$	ω_L	ω_R
a_L	λ_i	0
a_R	0	$1 - \lambda_i$

where $i \in \{DM, A\}$ and $0 < \lambda_A < \lambda_{DM} < 1$.[10] In particular, at some belief $q = Pr(\omega_R)$, player i has a higher expected payoff from a_R than from a_L if and only if

$$Eu_i(a_R) = q(1 - \lambda_i) \geq (1 - q)\lambda_i = Eu_i(a_L)$$

$$\Leftrightarrow$$

$$q \geq \lambda_i. \tag{1}$$

If $q < \lambda_A$ then both prefer a_L, if $\lambda_A < q < \lambda_{DM}$ the DM prefers a_L while the agent prefers a_R and if $\lambda_{DM} < q$ both prefer a_R (see Figure 1).

Under *observable testing*, in each period $t \in \{1, \ldots, T\}$, the agent chooses to test ($\tau = 1$) or to stop testing ($\tau = 0$). If the agent tests, a signal $s^t \in \{s_L, s_R\}$ is drawn, i.i.d. conditional on the state, where

$$Pr(s_i|\omega_i) = p, \tag{2}$$

with $\frac{1}{2} < p < 1$ and $i \in \{L, R\}$. If the agent stops, he cannot test again. Let s be a complete list of signal realizations, i.e. the list of realizations if the agent runs all T tests,

[8] The finite horizon represents the agent's resource constraints and ensures that players never fully learn. A finite horizon can be interpreted as a limiting case of an infinite horizon with a convex cost of testing.

[9] Consider a general payoff function for $i = A, DM$:

$u_i(a, \omega)$	ω_L	ω_R
a_L	α_i	β_i
a_R	γ_i	δ_i

where $\alpha_i, \beta_i, \gamma_i, \delta_i \in \mathbb{R}$ and $\alpha_i > \gamma_i$ and $\delta_i > \beta_i$. Then my results apply if and only if $\lambda_i = \frac{\alpha_i - \gamma_i}{(\alpha_i - \gamma_i) + (\delta_i - \beta_i)}$.

[10] If I instead assumed $0 < \lambda_{DM} < \lambda_A < 1$, the results would be a mirror image of the results presented: the labels of states $\omega \in \{\omega_L, \omega_R\}$ and actions $a \in \{a_L, a_R\}$ would have to be switched.

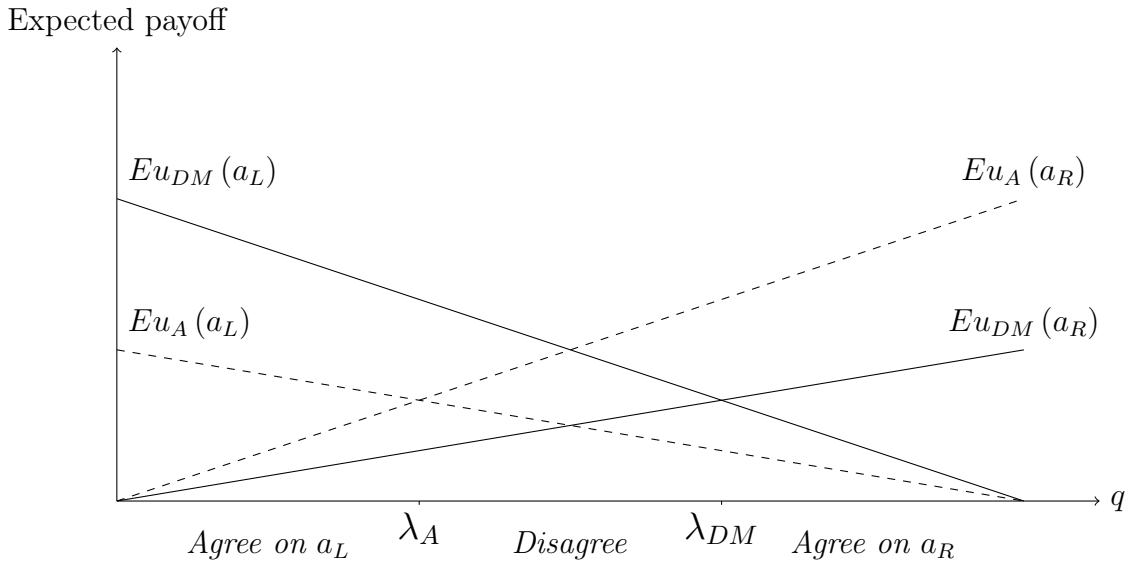

Figure 1: This figure illustrates each players' preferred action a at a given belief $q = Pr(\omega_R)$.

and let S be the set of such lists. Let $h_t = (s^1, \ldots, s^t)$ be the signal history at the end of period t, where s^k is the k^{th} element of s if the agent tests in period k and $s_\emptyset$ otherwise for $k \in \{1, \ldots, t\}$. Let H_t be the set of such histories h_t for $t \in \{1, \ldots, T\}$. A strategy for the agent is $\sigma_A : \cup_{t=1}^{T} H_{t-1} \to \{0, 1\}$. It selects action τ conditional on signal history h_{t-1} with $s_\emptyset \notin h_{t-1}$ for $t \in \{1, \ldots, T\}$. A strategy for the DM is $\sigma_{DM} : H_T \to \{a_L, a_R\}$. It selects action a conditional on final history h_T.

Under *hidden testing*, the agent's choice of τ and the realization of the signal are his private information. At the end of period T, the agent sends a message $m \in \mathcal{M}$. The message space is given by

$$\mathcal{M} \equiv \left\{ \left\{ s^1, ..., s^T \right\} | s^t \in \{s_L, s_R\}, t \in \{1, \ldots, T\} \right\}. \tag{3}$$

The set of feasible messages is $M(h_T) \subset \mathcal{M}$. $M(h_T)$ is the power set of the set that contains the same number of s_R and s_L realizations as history h_T, e.g. if $T = 2$ and $h_T = (s_R, s_L)$ then the set of feasible messages is given by $M((s_R, s_L)) = \{\emptyset, \{s_R\}, \{s_L\}, \{s_R, s_L\}\}$.[11] This

[11] This assumptions is made to capture a setting in which the DM can only make limited inferences from the calendar time of reports, as in Felgenhauer and Schulte [2014]. For example, this could be because both the time at which the agent first has the possibility to test or the total number of tests feasible in any interval of time is private information to the agent and cannot be verifiably disclosed by him.

means that the agent can report any subset of signals he has collected, but he cannot forge signals. The agent has a testing strategy, $\sigma_A^T : \cup_{t=1}^T H_{t-1} \to \{0,1\}$, and a reporting strategy. A reporting strategy for the agent is $\sigma_A^M : H_T \to \mathcal{M}$. It selects a message $m \in M(h_T)$ conditional on history h_T. A strategy for the DM is $\sigma_{DM} : \mathcal{M} \to \{a_L, a_R\}$. It selects action a conditional on message m.

In either game, the solution concept is a PBE in pure strategies and I assume that the DM chooses a_R if indifferent between a_L and a_R.[12] I impose the additional requirement that, for any signal history off the equilibrium path, players update their beliefs about the state according to Bayes' rule using the test accuracy given by (2).[13]

3 Observable Testing

When testing is observable, both players update their beliefs about the state using all past signal realizations and, hence, both have the same posterior belief at any point in time. The agent can influence the DM's belief about the state via controlling the flow of signals, that is he can strategically choose to stop testing at certain posterior beliefs but continue testing at others. However, when the agent stops testing, he also suffers the downside that he does not learn more himself. In this section, I will show that the agent may stop testing even if further tests could be beneficial to both players.

Let $Q(q_0)$ be the support of beliefs, that is the discrete set of posterior beliefs that are formed according to Bayes' rule given prior q_0 and given some signal history h_T as $T \to \infty$.[14] Let q_t be the belief in period t. One signal realization can lead the current belief to move either to the next highest or the next lowest belief in $Q(q_0)$:[15]

[12] The qualitative insights are unaffected by restricting attention to pure strategies.

[13] With this assumption, the set of equilibrium outcomes is the same as when the equilibrium concept is a sequential equilibrium.

[14] Beliefs move along a "grid" specified by $Q(q_0)$ rather than along a line because evidence is assumed to be discrete. Discrete evidence makes the analysis of observable testing more cumbersome, but is needed to restrict the agent's scope of manipulation from selective disclosure when testing is hidden.

[15] Any pair of histories h_t and h_s for $t, s \in \{0, \ldots, T\}$ lead to the same $q \in Q(q_0)$ if and only if the difference between the number of s_R-realizations and s_L-realizations in h_t is the same as in h_s.

$$q_{t+1} = \begin{cases} \inf \{q \in Q(q_0) : q_t < q\} & \text{if } s = s_R, \\ \sup \{q \in Q(q_0) : q_t > q\} & \text{if } s = s_L. \end{cases} \tag{4}$$

Denote the smallest belief in the support $Q(q_0)$ at which the DM prefers a_R by

$$\overline{q} \equiv \inf \{q \in Q(q_0) : \lambda_{DM} \leq q\}, \tag{5}$$

and the largest belief in the support $Q(q_0)$ at which the DM prefers a_L by

$$\underline{q} \equiv \sup \{q \in Q(q_0) : q < \lambda_{DM}\}. \tag{6}$$

I will impose two restrictions to make the analysis interesting. The first is that signal realizations can be pivotal to the DM's optimal choice. That is, there exists some signal list $s \in S$ and some testing strategy σ_A such that $q_T \geq \overline{q}$, and others such that $q_T \leq \underline{q}$. If this were not the case, the unique equilibrium outcome would be that the DM chooses the same action a irrespective of what the agent does. Second, I assume there is at least one belief at which players disagree, i.e. $\lambda_A < \underline{q} < \lambda_{DM}$. Without this assumption, players would agree on the optimal action a following any signal history and, hence, equilibrium outcomes would be the same as if there was no conflict of interest at all.

For both players, all payoff-relevant information contained in signal history h_t can be summarized by the posterior belief q_t for all $t \in \{0, \ldots, T\}$. Therefore, we can think of strategies as mapping posterior beliefs into action choices.

Proposition 1 *Under observable testing, the unique equilibrium strategies are as follows:*

1. *The DM chooses a_R if and only if $\overline{q} \leq q_T$.*

2. *For every $(\lambda_A, \lambda_{DM})$, there exists a critical number of remaining periods r' such that the agent strictly prefers to stop testing in period t if and only if $q_t = \overline{q}$ and $T - t < r'$.*

The DM's optimal strategy follows immediately from (1). The agent's payoff is unaffected by his decision to test or to stop if the remaining tests are not pivotal to the DM's final choice. Therefore, it is crucial to analyze if the agent tests at a belief just above or below the DM's threshold λ_{DM}. If the belief is just below the threshold, i.e. if $q = \underline{q}$, then an additional test must always be strictly beneficial for the agent. This is because, conditional on the test realization being pivotal to the DM's choice, the DM chooses the action a_R instead of a_L and, if the DM prefers a_R given the signal realization, then so does the agent. The same is not true at a belief just above the DM's threshold, i.e. $q = \overline{q}$. Conditional on the test being pivotal to the DM's choice, the DM chooses the action a_L, yet the agent may prefer action a_R. The continuation value of testing in this situation depends on how many tests the agent has left. In particular, if the agent only has a single test left then it is optimal for him to stop since he prefers a_R even following an s_L-realization.

Lemma 1 *Under observable testing, if $q_t = \overline{q}$ and $T - t \geq 3$, then there exists some λ_A such that the agent strictly prefers to stop testing, yet for some signal list $s \in S$ and some testing strategy σ_A the final belief satisfies $q_T < \lambda_A$.*

The agent may stop at belief $\overline{q}$ and the DM chooses a_R even though further tests could convince both players that a_L is the preferred action and, therefore, could yield a mutual benefit. Clearly, the agent would benefit from testing if the additional tests indeed turn out to lead both players to prefer a_L. However, when deciding whether to run further tests, the agent also considers the event that the additional signal realizations could cause the DM to choose a_L instead of a_R, even though the agent himself continues to prefer a_R. Therefore, for some realizations, the agent suffers a loss because the DM does not act in his interest, and this loss can outweigh the gain from supplying additional information. This logic behind Corollary 1 is illustrated in a simple set-up in the example below and we will return to this example in the subsequent section.

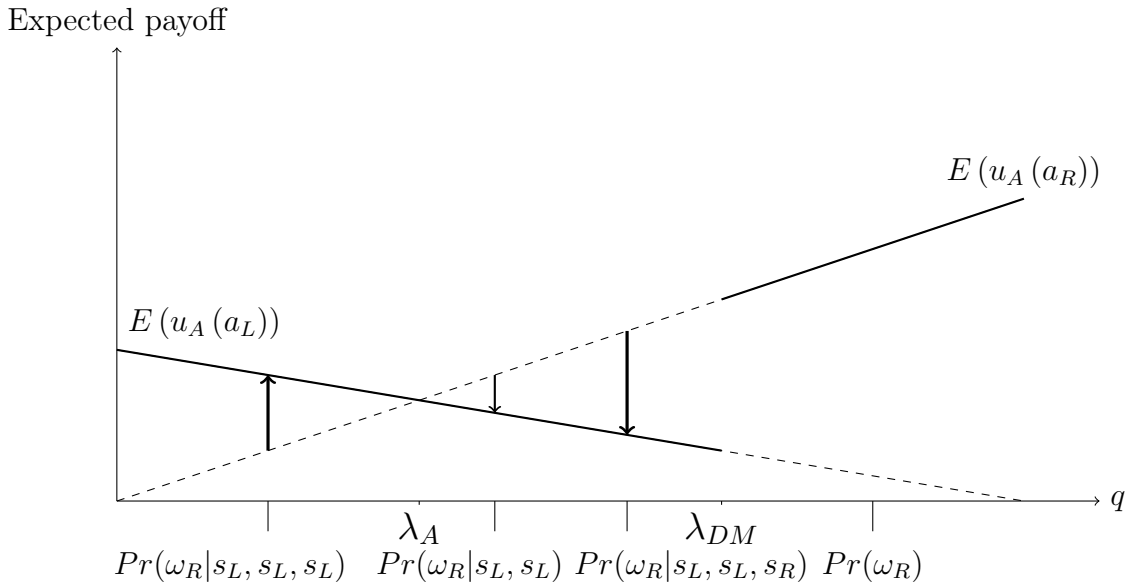

Figure 2: This figure illustrates the example in Section 3. The thick line indicates the agent's expected payoff if the DM follows his optimal strategy given belief q. If the agent does not test at all, the DM chooses $a = a_R$. If the agent tests and $h_3 = (s_L, s_L, s_L)$, the agent's expected payoff is higher than if he had not tested, as indicated by the upward-pointing arrow. If he tests and $h_3 = (s_L, s_L, s_R)$, the agent's expected payoff is lower than if he had not tested, as indicated by the downward-pointing arrow. The expected loss outweighs the expected gain from testing since the agent prefers a_L to a_R at $h_2 = (s_L, s_L)$. Therefore, the agent optimally does not test, although running tests could lead both players to learn that a_L is their preferred choice.

Example (see Figure 2). *Suppose* $T = 3$ *where*

$$Pr(\omega_R|s_L,s_L,s_L) < \lambda_A < Pr(\omega_R|s_L,s_L) < Pr(\omega_R|s_L) < \lambda_{DM} < Pr(\omega_R) = \overline{q}, \qquad (7)$$

that is, if all three tests indicate that the state is ω_L*, then both players prefer* a_L*, and if fewer than three tests are conducted and all indicate that the state is* ω_L*, then only the DM favors* a_L*, while at the prior belief, both players prefer* a_R*.*

In equilibrium, the agent does not test at all for the following reasons. Conditional on $s = (s_L, s_L, s_L)$*, the agent benefits from running all tests relative to not testing at all, because he prefers* a_L *and this is what the DM chooses if he observes* $h_3 = (s_L, s_L, s_L)$*. This is the only realization of signals for which the agent prefers* a_L *and, therefore, the only realization*

conditional on which he can strictly benefit from testing. Backward induction shows that the agent has no incentive to learn if realizations are indeed $s = (s_L, s_L, s_L)$. *At* $h_2 = (s_L, s_L)$, *the DM will choose* a_L *whatever the final signal realization is. Therefore, at the start of period 1, the agent anticipates that after the first two tests either* $h_2 = (s_L, s_L)$ *and the DM chooses* a_L *or, otherwise, the DM chooses* a_R. *Conditional on the event that running two tests is pivotal to the DM's choice, i.e. conditional on* $h_2 = (s_L, s_L)$, *the agent actually prefers the DM to stick to* a_R. *Therefore, the agent strictly prefers not to test.*

Lemma 1 does not apply when $T - t \leq 2$, i.e. when the agent has two or fewer tests remaining at $q = \overline{q}$. The reason is that for a mutual interest in learning to arise, the agent needs to prefer a_L when both tests yield s_L. But then there is no realization of the two tests for which the players disagree on the optimal action and, consequently, the agent only has an upside but no downside from running both tests.[16]

4 Hidden Testing

When testing is hidden, the agent has no strict benefit from stopping to test. What the DM learns about the state depends on what the agent chooses to reveal to him. In this section, I will show for which realizations the agent, through strategic disclosure, can lead the DM to choose an action that is not the DM's preferred action given the evidence. For the welfare comparisons of the next section, it is particularly interesting to analyze agent-preferred equilibria, i.e. equilibria in which the agent fully exploits the discretion he has under hidden testing to his advantage.

Since the order in which signals arrive does not affect the set of equilibrium outcomes under hidden testing, it is useful to track only the number of s_R-realizations contained in

[16] If the accuracy of the test is state-dependent and $Pr(s_L|\omega_L) > Pr(s_R|\omega_R)$, then Lemma 1 applies for $T - t \geq 2$ since it is possible that $\lambda_A < Pr(\omega_R|s_L, s_R) < \lambda_{DM} < Pr(\omega_R)$.

signal list $s \in S$. Define $N : S \to \mathbb{N}$ where

$$N(s) = \sum_{t \in \{1,\ldots,T\}} \mathbb{1}_{s^t = s_R}. \tag{8}$$

Let ϵ be an equilibrium under hidden testing and $\mathcal{E}$ be the set of all such equilibria.

Lemma 2 *Suppose testing is hidden.*

1. *For any equilibrium $\epsilon \in \mathcal{E}$, there exists a critical number $\overline{n}_\epsilon$ such that a_R is chosen given $s \in S$ in equilibrium ϵ if and only if $N(s) \geq \overline{n}_\epsilon$.*

2. *Any $\epsilon \in \mathcal{E}$ is Pareto-undominated.*

3. *In any equilibrium ϵ^* with*

$$\overline{n}_{\epsilon^*} = \min_{\epsilon \in \mathcal{E}} \overline{n}_\epsilon \tag{9}$$

the agent achieves the highest expected payoff and the DM achieves the lowest expected payoff across all $\epsilon \in \mathcal{E}$.

This structure arises because the agent's disclosure of signals is verifiable. The agent has no downside from running all tests and, if the agent optimally convinces the DM to choose a_R given some list of signals s, then he must also be able to convince the DM when the number of signals in favor of state ω_R increases.[17]

In what follows, we focus on agent-preferred equilibria, which yield the lowest equilibrium payoff to the DM, and characterize the critical number $\overline{n}_{\epsilon^*}$ which is a key feature of the agent's disclosure strategy σ_A^M in these equilibria.

Proposition 2 *In any agent-preferred equilibrium ϵ^* under hidden testing:*

1. *Let m_R be the message that the agent sends if $s = (s_R, \ldots, s_R)$. Then the agent sends this message m_R given $s \in S$ if and only if s satisfies $N(s) \geq \overline{n}_{\epsilon^*}$.*

[17] This uses the assumption that, for any signal history off the equilibrium path, players believe that signals were drawn from (2).

2. *Conditional on receiving message m_R, the DM holds posterior belief*

$$Pr\left(\omega_R|m_R\right) = Pr\left(\omega_R|N\left(s\right) \geq \overline{n}_{\epsilon^*}\right) \tag{10}$$

and chooses a_R.

3. *The critical number $\overline{n}_{\epsilon^*}$ satisfies*

$$\overline{n}_{\epsilon^*} = \min\left\{n \in \{0,\ldots,T\} : \lambda_{DM} \leq Pr\left(\omega_R|N\left(s\right) \geq n\right) \text{ and } \lambda_A \leq Pr\left(\omega_R|N\left(s\right) = n\right)\right\}. \tag{11}$$

In any agent-preferred equilibrium, there exist signal histories such that the DM chooses a_R following the agent's message, yet the DM would have preferred to choose a_L had he observed the signal history directly. The reason is that the agent pools signal histories with sufficiently many s_R-realizations such that the DM chooses a_R conditional on believing that the signal history was part of this pool. In particular, the agent pools histories by disclosing exactly $\overline{n}_{\epsilon^*}$ s_R-realizations for any h_T which contains $\overline{n}_{\epsilon^*}$ or strictly more s_R-realizations. Given this strategy, the DM is aware that the agent is hiding realizations, yet he strictly prefers to choose a_R conditional on exactly $\overline{n}_{\epsilon^*}$ s_R-realizations being disclosed. This threshold number $\overline{n}_{\epsilon^*}$ is as small as possible subject to the DM optimally choosing a_R conditional on inferring that $N\left(s\right) \geq \overline{n}_{\epsilon^*}$ and subject to the agent preferring a_R for each s such that $N\left(s\right) \geq \overline{n}_{\epsilon^*}$.

In such an equilibrium, the agent is indifferent between revealing exactly $\overline{n}_{\epsilon^*}$ s_R-realizations or all s_R-realizations if the DM chooses a_R for either message. This indifference is broken in favor of revealing exactly $\overline{n}_{\epsilon^*}$ s_R-realizations. One justification for this is that the agent could have an infinitesimal cost ε of running tests. Then the agent would strictly prefer to stop testing once he has exactly $\overline{n}_{\epsilon^*}$ s_R-realizations and not hide any realizations in equilibrium.

Corollary 1 *If $\lambda_{DM} \leq q_0$, then*

$$\overline{n}_{\epsilon^*} = \min\left\{n \in \{0,\ldots,T\} : \lambda_A \leq Pr\left(\omega_R|N\left(s\right) = n\right)\right\}. \tag{12}$$

When both players agree on a_R at the prior, the agent can strategically reveal realizations in such a way that the DM always chooses the agent's preferred choice. I will outline why this is the case in more detail below using the example introduced in Section 3.

Example *Consider again preferences described by (7). Then the following describes equilibrium play in an agent-preferred equilibrium. The agent runs all tests and his reporting strategy is*

$$m = \begin{cases} \{s_L, s_L, s_L\} & \text{for } h_T = (s_L, s_L, s_L) \\ \{s_R\} & \text{for } h_T \neq (s_L, s_L, s_L)\,. \end{cases} \tag{13}$$

The DM chooses a_R following $m = \{s_R\}$ and chooses a_L following $m = \{s_L, s_L, s_L\}$. The DM's response is optimal for the following reasons. Conditional on $m = \{s_L, s_L, s_L\}$, the DM correctly infers the signal history and chooses a_L because $Pr(\omega_R|s_L, s_L, s_L) < \lambda_{DM}$. Given that $m = \{s_L, s_L, s_L\}$ leads the DM to lower his belief that the state is ω_R, $m = \{s_R\}$ must raise his belief that the state is ω_R. Since the DM favors a_R at the prior, it must be that conditional on $m = \{s_R\}$, the DM chooses a_R. This means that the DM chooses a_R even if $h_T = (s_L, s_L, s_R)$, a history for which the DM would have preferred to choose a_L. Since the agent's most preferred decision is chosen for any realization of signals, the agent has no reason to deviate.

5 Hidden vs. Observable Testing

The findings of the previous sections allow us to characterize parameter combinations under which the agent acquires information pivotal to the DM's choice if and only if testing is hidden. When these conditions hold, hidden testing is always beneficial for the DM and can even be beneficial to both players. If the agent could commit to a disclosure strategy, the conditions would ensure that hidden testing would always benefit both players.

When testing is observable, the agent may stop at $q = \overline{q}$, as shown in Proposition 1. The agent will stop if he is too concerned that, in the end, the additional tests will lead the DM

to choose a_L when he himself still prefers a_R. This concern is stronger the larger the conflict of interest, i.e. the smaller λ_A. The following theorem shows that for intermediate values of λ_A, the agent always stops if $q = \overline{q}$ under observable testing, yet under hidden testing he is motivated to test and his messages help the DM to better match the action choice to the state. This leads to the following result:

Theorem 1 *For every λ_{DM} with $\lambda_{DM} \leq q_0$ there exist thresholds $\underline{\lambda}_A$ and $\overline{\lambda}_A$ such that the following two statements hold if and only if $\underline{\lambda}_A < \lambda_A < \overline{\lambda}_A$:*

- *In the unique equilibrium under observable testing, a_R is chosen for all $s \in S$.*
- *In any equilibrium under hidden testing, a_L is chosen for some $s \in S$ and, if a_L is chosen given $s \in S$, then the DM prefers a_L given s, i.e. $Pr(\omega_R|s) < \lambda_{DM}$.*

In particular, $\overline{\lambda}_A$ is chosen such that $\lambda_A < \overline{\lambda}_A$ if and only if the agent always stops at $q = \overline{q}$ under observable testing. In addition, $\underline{\lambda}_A$ is chosen such that $\underline{\lambda}_A < \lambda_A$ if and only if there exists some $s \in S$ such that the agent prefers a_L given s, that is, players have a mutual interest in learning. If $\underline{\lambda}_A < \lambda_A < \overline{\lambda}_A$, in the agent-preferred equilibrium under hidden testing, the agent tests because he enjoys the upside of the DM choosing a_L when it is mutually beneficial, and he does not suffer the downside of the DM choosing a_L when the agent prefers a_R, as shown by Corollary 1. The DM sometimes chooses a_R when he would have preferred a_L given the evidence collected, yet he at least learns if there was strong evidence in favor of the state being ω_L. If the DM chooses a_L given some signal list s' in the agent-preferred equilibrium, then he must also choose a_L given s' in any other equilibrium by Lemma 2.

Corollary 2 *If $\lambda_{DM} \leq q_0$ and $\underline{\lambda}_A < \lambda_A < \overline{\lambda}_A$, the following holds:*

- *In some equilibrium under hidden testing both players have a strictly higher payoff than in the unique equilibrium under observable testing.*

- *The DM has a strictly higher payoff in any equilibrium under hidden testing than in the unique equilibrium under observable testing.*

The key contribution of this paper is to show that both players can benefit from hidden testing. In an agent-preferred equilibrium, the implications of hidden testing are similar to those of a limited liability insurance issued to the agent by the DM. Since the agent can hide evidence if he his research leads the two players to disagree, the DM effectively commits to act in line with the agent's interest whatever his findings and the agent suffers no downside from testing. Although the DM does not necessarily take his preferred action given what the agent found out, assuring the agent that he will act in the agent's interest is what motivates the agent to test at all.

An additional motivation for focusing on agent-preferred equilibria under hidden testing in the comparison in Corollary 2 is that these equilibria give rise to the unique equilibrium payoff vector when the agent can commit to a disclosure strategy.[18]

Lemma 3 *Suppose the agent has the power to commit to a disclosure strategy σ_A^M when testing is hidden. Then Theorem 1 continues to apply and, if $\lambda_{DM} \leq q_0$ and $\underline{\lambda}_A < \lambda_A < \overline{\lambda}_A$, both the DM and the agent are strictly better off in any equilibrium under hidden testing than in the unique equilibrium under observable testing.*

6 Discussion and Relation with Existing Literature

The key contribution of this paper is to show that hidden testing can benefit both the DM and the agent. This has not been pointed out in the existing literature on the comparison between hidden and observable testing because it is usually assumed that the conflict of

[18]The assumption that the agent can commit to messages is made in the literature on Bayesian Persuasion, e.g. see Kamenica and Gentzkow [2011]. However, unlike in Bayesian Persuasion the agent does not have complete freedom in choosing the informativeness of the message.

interest changes monotonically with beliefs.[19] In closely related work, Henry [2009] and Felgenhauer and Loerke [2017] find that hidden testing is beneficial for the DM, but harmful for the agent across all Pareto-undominated equilibria. To compare their findings to those of this paper, I first give parameter combinations for which the DM is better off and the agent worse off in any equilibrium under hidden testing compared to the unique equilibrium under observable testing. Then I explain why the underlying effects arising in this situation are similar to the ones identified by their analyses, yet differ from those analyzed in Section 5. Finally, I will outline how the cost of testing plays a different role in Henry [2009]'s setting compared to the setting of this paper.

Consider $\lambda_A = 0$, which means the agent prefers a_R for all $s \in S$, and

$$Pr(\omega_R) < \lambda_{DM} < Pr(\omega_R|s_R), \tag{14}$$

that is, the DM prefers a_L at the prior, but prefers a_R if a single test is run and its realization is s_R. In equilibrium under observable testing, the agent tests and strictly prefers to stop if and only if the first signal realization is s_R. The agent does not value additional information and testing more has the downside that it could lead the DM to choose a_L instead. Hence, on path, the DM chooses a_R if and only if the first test yields s_R.

Under hidden testing, the agent runs both tests. Ideally, the agent would like the DM to choose a_R independent of the signal realizations, but if he reports the same message irrespective of his findings, the DM will choose a_L since he prefers a_L at the prior. For the agent, the next best candidate for equilibrium is one in which he pools histories with at least one s_R signal. However, conditional on knowing that at least one of two tests yielded s_R the DM still opts for a_L if

$$Pr(\omega_R|N(s) \geq 1) = \sum_{s \in \{(s_L,s_R),(s_R,s_L),(s_R,s_R)\}} Pr(\omega_R|s) < \lambda_{DM}. \tag{15}$$

[19] E.g. the agent has a preference over binary actions which is independent of the state or, when the state and action spaces are continuous, the agent prefers an action that differs by a constant amount from the DM's ideal action in any given state.

Assuming (15) holds, it must be that in equilibrium, the agent reveals all s_R signals and the DM chooses a_R if and only if both tests yielded s_R:

$$Pr\left(\omega_R | N\left(s\right) = 2\right) \geq \lambda_{DM}. \tag{16}$$

Clearly, in this two-period example, the DM gains from hidden testing. Under observable testing, the DM can only base his choice on the realization of a single test. By contrast, under hidden testing, he learns whether a realization in favor of state ω_R of the first test is backed up by the realization of the second test. The reason is that the agent is tempted to hide s_L-realizations and search for s_R that he can report instead. Since the DM anticipates that the agent is hiding adverse evidence, the DM requires to see additional s_R-realizations to select a_R when testing is hidden relative to when testing is observable. This leads the agent to test more than under observable testing but reduces the chances of approval. Hence, the agent is worse off under hidden testing.

The reasoning behind this example with $\lambda_A = 0$ is similar to the reasoning behind findings in the existing literature. Henry [2009] assumes that the agent's preferred action differs from the DM's preferred action by a constant in any given state. The agents chooses once and for all how much information to acquire at some cost. Under hidden testing, the agent invests more in information acquisition. This is because he benefits from having a larger sample to select from for any given reaction to his disclosure by the DM. However, in equilibrium, the DM anticipates how much evidence the agent has acquired and so the agent reveals all findings to convince the DM that he is not hiding adverse evidence.[20] In addition, Felgenhauer and Loerke [2017] assume that the agent prefers one of two possible actions independent of the state. The agent chooses how informative each test is and, under

[20]In Henry [2009]'s setting, the DM's action choice is continuous and any change in the DM's belief influences his optimal action choice and the agent's payoff. In my setting, the DM takes the same action for a range of beliefs, which means that the agent is indifferent between disclosures that lead to any one of these beliefs. If there was a small cost of testing, the agent would stop as soon as he can disclose enough to convince the DM to choose the agent's preferred action and the DM would not know how much evidence the agent has acquired in equilibrium.

hidden testing, runs a more informative test to make it credible that he has no interest in running any further test following an adverse finding.

The agent's motivation for running additional tests is different from the one outlined in Section 5. One approach to highlight this difference is to assume that the agent has the power to commit to which signals he discloses for a given set of findings. In the example above with $\lambda_A = 0$, the DM would no longer be strictly better off and the agent would no longer be strictly worse off under hidden testing than under observable testing if the agent had such commitment power. In particular, with commitment power, the agent could always commit to full disclosure and then the equilibrium outcome and equilibrium payoffs under hidden testing would be the same as under observable testing. However, as shown by Lemma 3, even with such commitment, Theorem 1 continues to apply and both players strictly benefit from hidden testing.

In addition, the cost of testing plays a different role in Henry [2009] than in the analysis of Section 5. Henry [2009] assumes that the agent incurs a cost for information acquisition. When the agent chooses the quantity of information publicly, he weighs up this cost against the benefit he derives from better informed action choices. When the quantity is chosen in private, the agent faces the same trade-off with the additional benefit that more information acquisition expands his scope for manipulation. By contrast, in the analysis of Section 5, under observable testing, the agent may stop acquiring information despite the fact that an additional test is costless. The reason for stopping is that more information can lead to a less preferred action choice from the agent's perspective. Under hidden testing, selective disclosure reduces the chances that the DM's action choice inflicts a loss on the agent which is why the agent runs more tests.

7 Further Examples

The insights of this paper also apply in the context of scientific research. While a researcher (agent) cares about learning the truth, he is also under pressure to publish and therefore may be less averse to accepting a false hypothesis than a journal (DM). The fact that the results of many scientific studies could not be replicated has strengthened demands for pre-analysis plans. However, with pre-analysis plans, once a researcher has a result that he could get published as is, he may avoid posing follow-up questions or run robustness checks. On the one hand, he would like to further investigate whether or not his findings hold up because he is interested in investigating the truth. On the other hand, he does not want to risk that additional findings cast some doubt on existing conclusions and lead the journal to reject the publication altogether when the scientist would have preferred to have it published.

In a political context, voters (DM) may be more reluctant to accept a new policy than the government (agent). Transparency in government allows voters to hold elected officials accountable for how they take decision based on evidence. However, with transparency, when the existing evidence leads voters to support the policy, the government may be deterred from collecting further evidence because this may cause voters to withdraw their support when the government would still like to see the policy implemented. By contrast, if the government does not have to disclose all the evidence, voters may benefit because the government would gather more evidence and dismiss the policy if it learns that it is very likely to be ineffective.

8 Conclusion

In many situations, a DM has to rely on an agent to become better informed about certain aspects on which her optimal choice depends. Since the DM and the agent's interests are not always aligned, a natural question is whether or not the decision maker should monitor more closely how the agent goes about collecting information. When information acquisition is public, the DM ensures that he can always act optimally given the evidence acquired,

whereas when information acquisition is private, the agent can manipulate the DM's choice through strategically revealing evidence. However, when information acquisition is public, the agent may stop acquiring evidence sooner than when acquisition is private and, therefore, the DM's decision is based on weaker evidence.

This paper shows that making the collection of information transparent can have adverse consequences for both the DM and the agent. When testing is observable, the agent may stop collecting evidence although further tests result in a superior choice for both players. This is because the agent fears that additional evidence may convince the DM to switch to an alternative action, yet be insufficient for the agent to prefer an alternative action. By contrast, under hidden testing, the agent can reduce the risk that the DM does not act in line with his interests by revealing evidence strategically. Therefore, the agent has an incentive to explore whether or not additional evidence leads both to agree on a different action choice. Although the DM does not necessarily act optimally given the evidence the agent has found, the DM benefits from learning whether or not the evidence was sufficient to convince the agent of an alternative action.

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

A Appendix

A.1 Proof of Proposition 1

Part 1) follows from (1) and (5). For Part 2), the first step is to show that the agent never has a strict benefit from stopping in period t if $q_t < \lambda_{DM}$. Denote his expected continuation value of stopping in period t at belief q by $V_t^0(q)$. If the agent stops at t, then $q_T = q_t$. Given the DM's strategy and $q_t < \lambda_{DM}$,

$$V_t^0(q_t) = q_t(1 - \lambda_A). \tag{17}$$

Denote the agent's expected continuation value when he tests in period t at belief q and then follows continuation strategy σ by $V_t^1(q, \sigma)$. Let σ^1 be the continuation strategy of testing in all remaining periods. Given the DM's strategy,

$$V_t^1\left(q_t, \sigma^1\right) = (1 - q_t) Pr\left\{q_T < \lambda_{DM} | \omega_L; \sigma^1\right\} \lambda_A + q_t Pr\left\{q_T \geq \lambda_{DM} | \omega_R; \sigma^1\right\} (1 - \lambda_A).$$

Then stopping cannot be optimal:

$$\begin{aligned} & V_t^1\left(q_t, \sigma^1\right) - V_t^0(q_t) \geq 0 \\ \Leftrightarrow (1 - q_t) Pr\left\{q_T < \lambda_{DM} | \omega_L; \sigma^1\right\} \lambda_A - q_t Pr\left\{q_T < \lambda_{DM} | \omega_R; \sigma^1\right\} (1 - \lambda_A) & \geq 0 \\ \Leftrightarrow Pr\left\{\omega_R | q_T < \lambda_{DM}; \sigma^1\right\} & \leq \lambda_A, \end{aligned}$$

and the inequality is satisfied since

$$\lambda_A > \lambda_{DM} = Pr\{\omega_R | q_T = \lambda_{DM}\} > Pr\left\{\omega_R | q_T < \lambda_{DM}; \sigma^1\right\}. \tag{18}$$

Consider $q_t > \overline{q}$. Let σ' be the continuation strategy of stopping in any period $t + n$ if

and only if $q_{t+n} = \overline{q}$. Then the agent has no strict benefit from stopping:

$$\begin{aligned} V_t^1(q_t, \sigma') &= (1 - q_t) \, Pr\left\{q_T < \lambda_{DM} | \omega_L; \sigma'\right\} \lambda_A + q_t Pr\left\{q_T \geq \lambda_{DM} | \omega_R; \sigma'\right\} (1 - \lambda_A) \\ &= q_t (1 - \lambda_A) = V_t^0(q_t). \end{aligned}$$

Consider $q_t = \overline{q}$. The agent strictly prefers to stop at $q = \overline{q}$ with only one test remaining, since

$$\begin{aligned} & V_T^1(\overline{q}, \sigma') - V_T^0(\overline{q}) < 0 \\ \Leftrightarrow & (1 - \overline{q}) \, p\lambda_A - \overline{q}(1 - p)(1 - \lambda_A) < 0 \\ \Leftrightarrow & \lambda_A < \frac{\overline{q}(1 - p)}{\overline{q}(1 - p) + (1 - \overline{q}) \, p} \equiv \underline{q}. \end{aligned}$$

and the last inequality holds by the assumption that players disagree for at least some realization, i.e. $\lambda_A < \underline{q} < \lambda_{DM}$.

The continuation value of testing $V_t^1(\overline{q}, \sigma')$ decreases with t. Let t' be the earliest period at which the agent has a strictly higher continuation value of stopping than of testing and then following strategy σ':

$$t' \equiv \inf\left\{t \in \mathbb{N} | V_t^1(\overline{q}, \sigma') - V_t^0(\overline{q}) < 0\right\}, \tag{19}$$

where

$$\begin{aligned} & V_t^1(\overline{q}, \sigma') - V_t^0(\overline{q}) < 0 \\ \Leftrightarrow (1 - \overline{q}) \, Pr\left\{q_T < \lambda_{DM} | \omega_L; \sigma'\right\} \lambda_A - \overline{q} Pr\left\{q_T < \lambda_{DM} | \omega_R; \sigma'\right\} (1 - \lambda_A) < 0 & \\ \Leftrightarrow \frac{\overline{q}}{1 - \overline{q}} \frac{Pr\left\{q_T < \lambda_{DM} | \omega_R; \sigma'\right\}}{Pr\left\{q_T < \lambda_{DM} | \omega_L; \sigma'\right\}} > \frac{\lambda_A}{1 - \lambda_A}. & \end{aligned} \tag{20}$$

Calculating $Pr\left\{q_T < \lambda_{DM} | \omega_i; \sigma'\right\}$ corresponds to calculating the probability of the event

that $q_{t+n} < \overline{q}$ for all $n \in \{1, \ldots, T-t\}$. In state ω_i for $i \in \{L, R\}$,

$$Pr\left\{q_T < \lambda_{DM} | \omega_i; \sigma^{'}\right\} = 1 - \left[Pr\left(s_R|\omega_i\right) + R\left(T-t\right)\right], \tag{21}$$

where $R : \{0, \ldots, T\} \to [0,1]$ is given by

$$\begin{aligned} R\left(T-t\right) &\equiv Pr\left(s_L|\omega_i\right) Pr\left(s_R|\omega_i\right) + \mathbb{1}_{T-t\geq 4} \sum_{k=1,..,\text{int}\left[\frac{T-t-2}{2}\right]} \frac{1}{k+1}\binom{2k}{k} \left(Pr\left(s_L|\omega_i\right)\right)^k \left(Pr\left(s_R|\omega_i\right)\right)^k \\ &= p\left(1-p\right) + \mathbb{1}_{T-t\geq 4} \sum_{k=1,..,\text{int}\left[\frac{T-t-2}{2}\right]} \frac{1}{k+1}\binom{2k}{k} p^k \left(1-p\right)^k, \end{aligned} \tag{22}$$

and $\text{int}\left[\frac{T-t-2}{2}\right]$ denotes the largest integer not strictly greater than $\frac{T-t-2}{2}$. To construct (21), first note that the event that $q_{t+n} < \overline{q}$ for all $n \in \{1, \ldots, T-t\}$ does not occur if and only if one of the following occurs: either $s_t = s_R$ (which implies $q_{t+1} > \overline{q}$), or $s_t = s_L$ and $q_{t+n} = \overline{q}$ for some $n \in \{2, \ldots, T-t\}$. The probability that $s_t = s_L$ and $q_{t+n} = \overline{q}$ for some $n \in \{2, \ldots, T-t\}$ is given by in (22). To construct $R\left(T-t\right)$, note that the event that $q = \overline{q}$ for the first time in period $t+n$ with $n \in \{2, \ldots, T-t\}$ can occur only if n is even. In particular, the event that $s_t = s_L$ and $q = \overline{q}$ for the first time in period $t+2k$ with $k \in \left\{1, .., \text{int}\left[\frac{T-t-2}{2}\right]\right\}$ occurs if and only if i) $s_t = s_L$ and $s_{t+2k} = s_R$ and ii) the intermediate $2k-2$ signals are equally split between s_R- and s_L-realizations and iii) the intermediate $2k-2$ signals arrive in an order such that in any period $t+n$ where $n \in \{1, ..., T-t-1\}$ the history h_{t+n} contains strictly more s_L- than s_R-realizations. By the properties of Dyck words, $\frac{1}{k+1}\binom{2k}{k}$ is the number of possible signal lists of length $2k$ such that (if the agent tests in each period) any history h_j with $j \leq 2k$ contains weakly more s_L- than s_R-realizations.

Hence, given (19), (20) and (21),

$$t' \equiv \inf\left\{t \in \mathbb{N} | \frac{\lambda_A}{1-\lambda_A} < \frac{\overline{q}}{1-\overline{q}} \frac{1-p-R\left(T-t\right)}{p-R\left(T-t\right)}\right\} \tag{23}$$

and the critical number of remaining periods is

$$r' \equiv \sup \left\{ r \in \mathbb{N} | \frac{\lambda_A}{1-\lambda_A} < \frac{\overline{q}}{1-\overline{q}} \frac{1-p-R(r)}{p-R(r)} \right\}. \tag{24}$$

A.2 Proof of Lemma 1

Let r be the number of tests remaining, i.e. $r = T - t$. Choose λ_A such that

$$\frac{\overline{q}(1-p)^r}{\overline{q}(1-p)^r + (1-\overline{q})p^r} < \lambda_A < \frac{\overline{q}(1-p)^{r-1}}{\overline{q}(1-p)^{r-1} + (1-\overline{q})p^{r-1}}. \tag{25}$$

If the testing strategy is to run all remaining r tests and $s^n = s_L$ for all $n \in \{t+1, \ldots, T\}$, then

$$q_T = \frac{\overline{q}(1-p)^r}{\overline{q}(1-p)^r + (1-\overline{q})p^r}, \tag{26}$$

and, hence, $q_T < \lambda_A$ for some testing strategy σ_A and some signal list s.

What remains to be shown is that the agent strictly prefers to stop testing at $q_t = \overline{q}$. The agent gains from testing if and only if the DM chooses a_L given q_T and the agent prefers a_L at q_T. Hence, given (25), the agent gains from testing at all if and only if q_T satisfies (26). For q_T to satisfy (26), there must exist some $n \in \{1, \ldots, r\}$ such that in period $t+n$ the belief satisfies

$$q_{t+n} = \frac{\overline{q}(1-p)^{r-1}}{\overline{q}(1-p)^{r-1} + (1-\overline{q})p^{r-1}}. \tag{27}$$

For (27) to hold, it must be that $n \geq r-1$ since $q_t = \overline{q}$. Hence, if the agent were to reach the belief in (27) in some period, he can have at most one test remaining. Conditional on q_{t+n} satisfying (27) and $n \geq r-1$, the highest possible belief in period T is

$$q_T = \frac{\overline{q}(1-p)^{r-2}}{\overline{q}(1-p)^{r-2} + (1-\overline{q})p^{r-2}}. \tag{28}$$

Given $r \geq 3$, it must be that

$$\frac{\overline{q}\left(1-p\right)^{r-2}}{\overline{q}\left(1-p\right)^{r-2}+\left(1-\overline{q}\right)p^{r-2}} < \overline{q}. \tag{29}$$

This shows that, conditional on the agent reaching the belief given by (27) in some period, the DM chooses a_L irrespective of the remaining realizations, i.e. $q_T < \lambda_{DM}$ with probability 1. Hence, the continuation value of testing in period t is the same whether the continuation strategy is to always test or whether the continuation strategy (denoted by $\widetilde{\sigma}$) is to test and stop if and only if (27) is satisfied for some $n \in \{1, \ldots, r\}$. Given continuation strategy $\widetilde{\sigma}$, the agent prefers a_R for any list of signal realizations, i.e. $Pr\left(\lambda_A < q_T | \omega_i, \widetilde{\sigma}\right) = 1$ for $i \in \{L, R\}$. Therefore, there is no upside to testing in period t. However, there is a downside to testing in period t. Testing leads the DM to choose a_L with strictly positive probability since $Pr\left(q_T < \lambda_{DM} | \omega_i, \widetilde{\sigma}\right) > 0$, while stopping results in the DM choosing a_R with probability 1.

A.3 Proof of Lemma 2

1.) Suppose the agent runs all T tests.[21] If a_R is chosen given s, then the agent must prefer a_R given s, i.e. $\lambda_A \leq Pr\left(\omega_R | s\right)$. If this were not the case, the agent would have a profitable deviation to send a message that contains all realizations in s. This deviation is profitable because $Pr\left(\omega_R | s\right) < \lambda_A$ and $\lambda_A < \lambda_{DM}$ imply that $Pr\left(\omega_R | s\right) < \lambda_{DM}$, which means the DM chooses a_R. In addition, if $\lambda_A \leq Pr\left(\omega_R | s\right)$ then $\lambda_A \leq Pr\left(\omega_R | s'\right)$ for any s' and s where $N\left(s'\right) \geq N\left(s\right)$.

The final step is to show that, if there exists a message $m \in M\left(s\right)$ such that $\sigma_{DM}\left(m\right) = a_R$, there also exists a message $m' \in M\left(s'\right)$ such that $\sigma_{DM}\left(m'\right) = a_R$ where $N\left(s'\right) \geq N\left(s\right)$. If $m \in M\left(s'\right)$ then choose $m' = m$. If $m \notin M\left(s'\right)$, then it must be that $N\left(s'\right) > N\left(s\right)$. If $N\left(s'\right) > N\left(s\right)$ and m' contains all signals in s', then consistent beliefs by the DM must

[21] The agent never has a strict downside from running all tests given the DM's strategy.

satisfy $Pr(\omega_R|m') = Pr(\omega_R|s') \geq Pr(\omega_R|m)$ given the requirement that players update their beliefs about the state according to Bayes' rule using the test accuracy given by (2) on or off the equilibrium path. In particular, given m' contains all signals in s', the DM must infer that $s = s'$ and, hence, that $N(s') \geq N(s)$ for any s such that $m \in M(s)$. If $\lambda_{DM} < Pr(\omega_R|m)$ then also $\lambda_{DM} < Pr(\omega_R|s')$ and, hence, $\sigma_{DM}(m') = a_R$.

2.) By the proof of Part 1., a_R cannot be chosen given s if both players prefer a_L given s, i.e. if $\lambda_A > Pr(\omega_R|s)$. In addition, a_L cannot be chosen given s in equilibrium if both players prefer a_R given s, i.e. if $\lambda_{DM} < Pr(\omega_R|s)$, because, if this were the case, the agent would have a profitable deviation to disclose all signal realizations which would lead the DM to choose a_R. Finally, for all s such that $\lambda_A < Pr(\omega_R|s) < \lambda_{DM}$, any change in the action a given s must strictly decrease at least one player's payoff.

3.) This follows from Part 1. and 2. given $\lambda_A < \lambda_{DM}$.

A.4 Proof of Proposition 2

The first step is to establish that Proposition 2 is consistent with equilibrium. It is weakly optimal for the agent to run all tests given σ_{DM}. If the agent's strategy σ_A^M is to send some $m = m_R$ given s if and only if $N(s) \geq \overline{n}_{\epsilon^*}$, then (10) gives the DM's consistent beliefs given m_R updated according to Bayes' rule. It is optimal for the DM to choose a_R given his beliefs since (11) implies $\lambda_{DM} \leq Pr(\omega_R|N(s) \geq \overline{n}_{\epsilon^*})$. Given that the DM chooses a_R given m_R, the agent optimally sends m_R given s only if he himself prefers a_R given s, i.e. only if $\lambda_A \leq Pr(\omega_R|s)$. This is the case since (11) implies $\lambda_A \leq Pr(\omega_R|N(s) = \overline{n}_{\epsilon^*})$ and $Pr(\omega_R|N(s) = \overline{n}_{\epsilon^*}) \leq Pr(\omega_R|N(s) \geq \overline{n}_{\epsilon^*})$. Finally, for the agent to choose m_R given s only if $N(s) \geq \overline{n}_{\epsilon^*}$, choose m_R such that $m_R \in M(s)$ if and only if $N(s) \geq \overline{n}_{\epsilon^*}$.[22]

The next step is to establish that Proposition 2 is the partial characterization of an agent-preferred equilibrium. There always exists an equilibrium ϵ in which the agent reports all signal realizations in s, the DM correctly infers s given the message and chooses a_R if

[22] E.g. let m_R contain $\overline{n}_{\epsilon^*}$ s_R-realizations and no s_L-realizations.

and only if $\lambda_{DM} \leq Pr(\omega_R|s)$. Let S_{DM} be the set of signal lists for which the DM prefers a_R, i.e.

$$S_{DM} = \{s \in S : \lambda_{DM} \leq Pr(\omega_R|s)\}. \tag{30}$$

Then there also exists an equilibrium ϵ' with the same payoff vector as ϵ in which the agent reports m_R if and only if $s \in S_{DM}$, the DM's beliefs conditional on m_R are

$$Pr(\omega_R|m_R) = \sum_{s \in S_{DM}} Pr(\omega_R|s) \tag{31}$$

and the DM chooses a_R if and only if $m = m_R$. That is, the agent pools all signal lists for which the DM prefers a_R.

Is there an equilibrium in which the agent does better? Potentially, the agent has the option to pool signal lists for which the DM prefers a_R with signal lists for which the DM prefers a_L. Suppose in an agent-preferred equilibrium, the agent sends m_R if and only if $s \in S^*$, where $S^* \subseteq S$. For this to be an equilibrium, it must be that the DM prefers choosing a_R following m_R given consistent beliefs, i.e.

$$\lambda_{DM} \leq Pr(\omega_R|m_R) = \sum_{s \in S^*} Pr(\omega_R|s), \tag{32}$$

and for all $s \in S^*$ the agent must prefer a_R, i.e.

$$\lambda_A \leq Pr(\omega_R|s). \tag{33}$$

Since each $s \in S^*$ such that $\lambda_{DM} \leq Pr(\omega_R|s)$ relaxes constraint (32), it must be that $S_{DM} \subset S^*$. Furthermore, $Pr(\omega_R|s)$ is higher, the higher $N(s)$. In addition, the agent's payoff from a_R given s is higher, the higher $N(s)$. Lastly, if $s \in S^*$ and $N(s) = n$, then all s' with $N(s') = n$ must also satisfy $s' \in S^*$ because otherwise the agent has a profitable deviation to send m_R for s' since $M(s') = M(s)$ and (33) for all $s \in S^*$. Therefore, there

must exist some cut-off $\overline{n}$ such that

$$S^* = \{s \in S : N(s) \geq \overline{n}\} \tag{34}$$

and this cut-off $\overline{n}$ must be as low as possible subject to (32) and (33) being satisfied.

A.5 Proof of Corollary 1

Suppose that the agent were to send m_R given s if and only if $s \in S^*$ where

$$S^* = \{s \in S : \lambda_A \leq Pr(\omega_R|s)\}. \tag{35}$$

Then this implies that the DM optimally chooses a_R given m_R, i.e. $\lambda_{DM} \leq Pr(\omega_R|m_R)$, since with consistent beliefs

$$\lambda_{DM} \leq q_0 = \sum_{s \in S} Pr(\omega_R|s) \leq \sum_{s \in S^*} Pr(\omega_R|s) = Pr(\omega_R|m_R). \tag{36}$$

Hence, (11) reduces to (12).

A.6 Proof of Theorem 1

First, find $\overline{\lambda}_A$ such that, $\lambda_A < \overline{\lambda}_A$ if and only if in the unique equilibrium under observable testing, the DM chooses a_R for all $s \in S$. Under observable testing, the DM chooses a_R for all $s \in S$ if and only if $Pr(\lambda_{DM} \leq q_T) = 1$. By Proposition 1, the agent strictly prefers to stop only if $q = \overline{q}$. Find the smallest number of periods needed to reach belief $\overline{q}$ from q_0, denoted by k, by solving

$$\overline{q} = \frac{q_0 (1-p)^k}{q_0 (1-p)^k + (1-q_0) p^k}. \tag{37}$$

By the proof of Proposition 1, the agent strictly prefers to stop at $q = \overline{q}$ in period k if and only if

$$\frac{\overline{q}}{1-\overline{q}}\frac{1-p-R(T-k)}{p-R(T-k)} > \frac{\lambda_A}{1-\lambda_A}. \tag{38}$$

where $R(\cdot)$ is given by (22). If the agent stops at $q = \overline{q}$ when he has $T-k$ periods remaining, he must also stop when he has fewer periods remaining. Therefore, choose

$$\overline{\lambda}_A \equiv \sup\left\{\lambda_A \in \mathbb{R} : \frac{\lambda_A}{1-\lambda_A} < \frac{\overline{q}}{1-\overline{q}}\frac{1-p-R(T-k)}{p-R(T-k)}\right\}. \tag{39}$$

Second, find $\underline{\lambda}_A$ such that, in any equilibrium under hidden testing, the DM chooses a_L for some $s \in S$. If the DM chooses a_L for some $s \in S$ in the agent-preferred equilibrium under hidden testing, then he will do so in any equilibrium under hidden testing given Part 1 of Lemma 2. Since $q_0 \leq \lambda_{DM}$, by Corollary 1, in any agent-preferred equilibrium, the DM chooses a_L given s if and only if $Pr(\omega_R|s) < \lambda_A$. Therefore, choose

$$\underline{\lambda}_A \equiv \frac{q_0(1-p)^T}{q_0(1-p)^T + (1-q_0)p^T} \tag{40}$$

such that if $\underline{\lambda}_A < \lambda_A$ then $Pr(\omega_R|s) < \lambda_A$ for $s = (s_L, ..., s_L)$. Since $\lambda_A < \lambda_{DM}$, if $\underline{\lambda}_A < \lambda_A$ then $Pr(\omega_R|s) < \lambda_{DM}$ for $s = (s_L, ..., s_L)$.

A.7 Proof of Corollary 2

The first statement follows by Theorem 1 and Corollary 1. The second statement follows directly from Theorem 1.

A.8 Proof of Lemma 3

Suppose the agent were to commit to the disclosure strategy of the agent-preferred equilibrium in the game without commitment power. Then the unique vector of equilibrium payoffs would be the same as in the agent-preferred equilibrium in the game without commitment

power. The agent has no reason to deviate and commit to another disclosure strategy because the action a chosen in equilibrium given s is his preferred action given s for every $s \in S$ by Corollary 1.

References

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.

M. Dahm, P. González, and N. Porteiro. Trials, tricks and transparency: How disclosure rules affect clinical knowledge. *Journal of Health Economics*, 28(6):1141 – 1153, 2009.

A. Di Tillio, M. Ottaviani, and P. N. Sorensen. Strategic Sample Selection. CEPR Discussion Papers 12202, C.E.P.R. Discussion Papers, Aug 2017a.

A. Di Tillio, M. Ottaviani, and P. N. Sørensen. Persuasion Bias in Science: Can Economics Help? *The Economic Journal*, 127(605):F266–F304, 2017b.

R. Dur and O. Swank. Producing and Manipulating Information. *The Economic Journal*, 115(500):185–199, 2005.

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.

T. Gall and Z. Maniadis. Evaluating solutions to the problem of false positives. *Research Policy*, 48(2):506 – 515, 2019.

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. *American Economic Review*, pages 911–55, March 2019.

M. Janssen. Generating and Suppressing Hard Evidence. *Unpublished Manuscript*, 2018.

T. Jefferson, MA. Jones, P Doshi, CB Del Mar, R Hama, MJ Thompson, EA Spencer, IJ Onakpoya, KR Mahtani, D Nunan, J Howick, and CJ Heneghan. Neuraminidase inhibitors for preventing and treating influenza in adults and children. *Cochrane Database of Systematic Reviews*, (4), 2014.

E. Kamenica and M. Gentzkow. Bayesian Persuasion. *American Economic Review*, 101(6): 2590–2615, October 2011.

S. Matthews and A. Postlewaite. Quality Testing and Disclosure. *The RAND Journal of Economics*, 16(3):328–340, 1985.

A. McClellan. Experimentation and Approval Mechanisms. *Unpublished working paper*, 2017.

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.

R. C. Rabin. "Suicide Data Incorrectly Reported in Drug Trials, Suit Claimed". *New York Times*, September 2017.

S. Shavell. Acquisition and Disclosure of Information Prior to Sale. *The RAND Journal of Economics*, 25(1):20–36, 1994.