CMS-EMS
Center for Mathematical Studies in Economics
And Management Science

Discussion Paper #1588

"Mediated truth"

Bruno Strulovici

Northwestern University

June 14, 2017

Mediated truth *

Bruno Strulovici
Northwestern University

– Very Preliminary –

June 14, 2017

Abstract

Many facts are learned through the intermediation of individuals with special access to information, such as law enforcement officers, employees with a security clearance, or experts with specific knowledge. This paper considers whether societies can learn about such facts when information is costly to acquire, cheap to manipulate, and produced sequentially. The answer is negative under an "asymptotic scarcity" condition pertaining to the amount of evidence available which distinguishes, for example, between reproducible scientific evidence and the evidence generated by a crime.

*Early versions of this project were presented at Microsoft Research New England (Summer 2015), Boston College, Washington University, Universitat Pompeu Fabra, Universitat Autònoma de Barcelona, Johns Hopkins University, University of Chicago, Cambridge University INET Economic Theory Conference, Penn State University, Cornell University, Arizona State University, Paul Milgrom's student seminar, Duke University, Northwestern University, and Bocconi's workshop on Advances in Information Economics. Part of this research was developed while I was visiting Microsoft Research New England and Harvard University whose hospitalities are gratefully acknowledged. This research is supported by an NSF Grant (No.1151410) and a fellowship from the Alfred P. Sloan Foundation.

1 Introduction

In order to function, societies and organizations must rely on individuals with specific access to information: criminal investigations are run by law-enforcement officers, political investigations are handled by special prosecutors, and scientific knowledge concerning policy-relevant factors such as the anthropogenic component of climate change and the effects of vaccination and tobacco on health, is also mediated by scientists with a special training and/or access to data. For the public at large as well as for policymakers and other economic agents, there is often no other way of learning about facts than through specific intermediaries, and even the "exogenous" signal often featured in economic models, perhaps an "accidental" discovery of evidence, must be reported by someone whose motives and discovery may potentially be questioned.

This mediated information is, moreover, often costly to acquire and often cheaper to fabricate or manipulate. For instance, it may be easier to coerce a false testimony or confession than to find an actual witness with reliable information, or to spread a rumor about someone rather than check what this person actually did, and it is easier to produce "significant" results when data can be fabricated or falsified. Sometimes, information manipulation may take more insidious forms, such as the selective or biased reporting of news by media outlets and the selective reporting of experimental results in scientific papers. Politicians, lawyers, and experts in various fields may also find it easier to produce arguments for or against a decision than to study seriously the actual merits of the decision.

In these and other contexts, incentivizing investigators raises the following dilemma: if incentives are too weak, investigators may prefer to shirk rather than acquire costly information. And if incentives are too highly powered, this creates a temptation for investigators to manipulate their findings in specific ways. For example, a researcher or expert witness paid by a corporation may have an incentive to report findings or opinions which align with the corporation's objectives. A prosecutor rewarded for achieving a high conviction rate may try to maximize this rate at the expense of the truth. A police department rewarded for a high clearance rate or investigating a particularly heinous or prominent crime may be pressured to "solve the case" at all cost, possibly discounting information that would undermine its main lead.

The natural answer to this dilemma is to compare the evidence produced by different investigators. For example, a judge rendering a verdict may be concerned that his opinion be corroborated by subsequent judicial opinions. A criminal investigator may face sanctions if his findings are contradicted by a higher-level investigation. In many applications, this corroboration process is sequential: at any point in time, there is one investigator in charge of the case, whose findings are disclosed at least to any future investigator of the case. Similarly, journalists and researchers can read articles published by their peers on a topic before producing their own article. There are instances in which the investigation process is parallel,[1] but in many applications the investigation process may be more accurately captured by a sequential description.

This paper introduces a model of sequential investigations with the following structure. The object of the investigation is a fact which may be learned through evidence uncovered by a costly process. In each period, an investigator is assigned to learn about the fact, and can choose between shirking, investigating the case seriously, or fabricating some evidence at a lower cost. The evidence produced is made available to subsequent investigators. An investigator's compensation can depend arbitrarily on all reports produced by the investigation, subject to a uniform bound on rewards and punishments.

The paper's main objective is to study whether there exist compensation schemes that induce learning: Given a prior distribution over the fact and the evidence generating process, can investigators' compensation be structured to induce adequate investigation and disclosure of evidence?

The answer, in turns out, depends on the evidence generating process. When evidence is reproducible, as is typically the case in mathematics and physics, it is possible to induce serious investigations until a certain level of confidence is gained about the underlying fact. For example, suppose that a sequence of mathematicians is asked whether a particular proof is correct, or a sequence experimental physicists is asked to check the value of a physical constant. There is no shortage of evidence for checking such facts: each investigator can generate her own piece of evidence. And investigators in each sequence can be incentivized to do so as long as there is enough uncertainty about the answer.[2]

When evidence is in limited supply, however, the result is negative. No matter how incentives are structured, it is impossible to make anyone work and all equilibria of the investigation process are uninformative. The limited supply assumption may be described as follows: as the body of real evidence already accumulated becomes arbitrarily large, it becomes arbitrarily unlikely that there exists additional evidence to uncover. The assumption does not preclude the possibility that some piece of evidence is indicative of more evidence to be found. But as more and more evidence

[1] A notable example concerns academic refereeing, although in this case the incentive structure is arguably rather remote from a Crémer-McLean mechanism.

[2] When confidence becomes very high, it becomes more profitable for investigators to herd with earlier findings rather than investigate the fact themselves. As long as there is sufficient uncertainty about the fact, however, outcome-based rewards can be used to incentivize investigators. The confidence threshold depends, among other things, on the cost of investigation relative to the maximal size of the rewards and punishment.

accumulates, the assumption implies that such this possibility vanishes.

To understand this result, the problem may be restated as follows: when is it possible to structure incentive so that investigators coordinate on the truth, despite the sequentiality of the investigation? Intuitively, to elicit the truth from an investigator, we must have available a signal that correlates with the truth above and beyond the information available to this investigator. As the investigation progresses, however, this signal becomes more difficult to acquire. Suppose that there exists an informative equilibrium. Then, either the evidence keeps accumulating, in which case the probability that further evidence exists becomes too small to warrant effort, or the process reaches a point at which no new evidence is reported for a long period of time, which is also indicative that there is no evidence left. Either way, investigators become increasingly pessimistic about the existence of additional evidence to be uncovered and stop working.

Thus the investigation process must reach an informational cascade of sort, in which all investigators' actions (whether they shirk or fabricate evidence) are no longer informative and purely based on previous investigators' effort. Unlike in the seminal herding models, however, this cascade ripples back to the early investigators, whose payoffs are completely dependent on future investigators' actions: if no investigator after round 11, say, ever seeks the truth, then investigator 10 has no incentive to seek it either, since other than through investigators 1–9 reports, which are publicly available, the truth will not affect 10's compensation.

There are two main challenges to establish this result. First, note that the unraveling argument is straightforward if one can establish that no investigator ever works after some fixed round *across all possible histories*. However, in any candidate equilibrium, the number of investigators who work seriously is history dependent and potentially unbounded (this is true even when evidence is reproducible prone to mistakes). When an investigator considers working, he understands that whether further investigations take place will depend on his report. For example, the investigator may be able to close the investigation by reporting that he found nothing, say, but keep it alive by claiming that he discovered new evidence.

Secondly, there is no obvious variable on which to build the "increased pessimism" argument. For example, some evidence produced by an investigator may suggest that past evidence was fabricated (e.g., if this evidence is incompatible with past evidence). Discovery of evidence, even taken at face value, does not necessarily make investigators believe that more evidence has been found, even in a stochastic sense. Likewise, if an investigator reports that he found nothing, this need not imply make it more likely that no evidence exists: for example, if this was the first round of investigation and it is common knowledge that there exist at least two pieces of evidence about the case, failure to find evidence reveals only that this investigator shirked or was unlucky.

2 Model

The object of the investigation is a fact, ω, lying in some measurable space Ω. For example, ω may represent whether a defendant is guilty, in which case Ω is binary, or the identity of the author of a crime, in which case Ω is the list of all suspects.

Associated with this fact is a collection E of pieces of evidence $(e_1, \ldots, e_M)$. In a criminal case, for instance, this represents all the evidence that the criminal has left behind. Ex ante, the collection E is random, including the number M of pieces. Let $\mathcal{E}$ denote the set of all possible evidence collections which may be generated, also assumed to be measurable, and $F(\cdot|\omega)$ denote the distribution of evidence over $\mathcal{E}$ conditional on each ω. For example, if ω represents whether a defendant is guilty and $\Omega = \{g, n\}$, e_m is incriminating if the marginal distribution of F with respect to e_m satisfies $F(e_m|g) > F(e_m|n)$.

In each round $i \geq 1$, a new individual is assigned to investigate the fact. i can choose between three actions: investigate the fact seriously at cost $c > 0$, shirk, or fabricate evidence at cost bounded by $d < c$.

For simplicity, fabrication and serious investigation are treated as mutually exclusive actions. However, the paper's main result holds as stated if one allows i to be fabricate evidence after investigating the fact, regardless of what he found, as explained in Section 6.

If i investigates the facts seriously (hereafter, "works") and there is some evidence and there is some real evidence left E'' to be uncovered, then i discovers some subset of this evidence with probability $\lambda_{E''}$. If all the evidence has been uncovered by past investigators, then a serious investigation reveal nothing (i.e., $\lambda_{\emptyset} = 0$).

The result of i's action is a collection m_i of evidence. This evidence, whether authentic or fabricated, is made available to all subsequent investigators.[3] The paper's main interest lies in the case where evidence is cheap to fabricate.

The incentives—punishments and rewards—provided to i are captured by a compensation function C_i, which may depend on the entire evidence sequence $m = (m_1, \ldots, m_j, \ldots)$. In particular, i may be punished if the evidence that he produced is incompatible with what is found by later investigators, or if he failed to produce evidence that is later uncovered by subsequent investigators. Alternatively, i may be rewarded if his finding seems novel compared to past investigators. The compensation function may represent a stream of rewards and punishment. For example, if i receives

[3]In applications, this evidence could be made public, as in the case of judicial opinions and published articles, or simply transferred to anyone inheriting the case, as in criminal investigations

a compensation $c_{i,j}(m_1, \ldots, m_j)$ at round $j \geq i$ and discounts future compensation according some discount factor $\delta < 1$, then

$$C_i(m_1, \ldots) = \sum_{j \geq i} \delta^{j-i} c_{i,j}(m_1, \ldots m_j)$$

The function C_i does not depend directly on the state ω. This assumption captures the idea that social planner (or public) in charge of administering the compensation does not get to observe the truth directly: its only way to assess i's efforts is through other investigators' reports, rather than through a public signal. We also assume that the punishments and rewards that can be given to any investigator are uniformly bounded and cannot depend significantly on evidence produced arbitrarily far in the future.

Given a realized stream $m = (m_1, \ldots,)$ of evidence, i's realized utility is

$$u_i(m) = C_i(m) - c_i$$

where $c_i = c$ is i sought evidence seriously and 0 otherwise. The equilibrium concept is the standard notion of Perfect Bayesian Equilibrium (PBE).

3 Reproducible evidence: An Example

Suppose that the state of the world is binary $\omega \in \{H, L\}$ with prior $P(\omega = H) = \hat{p}$ and the set $\mathcal{E}$ consists of all sequences of signals "H" and "L". If i investigates the fact, he receives an informative signal $\tilde{m}_i$ about ω such that $Pr(“H''|H) = Pr(“L''|L) = \pi > 1/2$. In this setting, all investigators discover some piece of evidence, so shirking can easily be discouraged by punishing any empty report. Thus suppose that i can either "work" (i.e., investigate the fact seriously) at cost $c > 0$ and report his evidence, i.e., $m_i = \tilde{m}_i$, or fabricate some evidence at zero cost, i.e., $m_i = a_i$ where $a_i \in \{“H'', “L'')$ is the evidence fabricated by i.

Given any equilibrium of this game, let for each $i \in \mathbb{N}$ γ_i denote the probability that i works, and let p_i denote the probability that $\omega = H$ given past information.

Proposition 1 *For any cutoffs $\underline{p}, \bar{p}$ such that $0 < \underline{p} < \hat{p} < \bar{p} < 1$, there exist large reward/punishment values P, R such that the following strategies constitute an equilibrium: $\gamma_i = 1$ as long as $p_i \in (\underline{p}, \bar{p})$ and $\gamma_i = 0$ afterwards. In particular, there exist informative equilibria.*

Proof. Given the symmetric signal structure, notice that the posterior p_i is entirely determined by the difference in the number of "H" and "L" signals as long as all agents $j < i$ work with probability 1. Therefore, the set of equilibrium posteriors forms a grid $\{q^k\}$ containing $\hat{p}$ and interrupted immediately outside of $(\underline{p}, \bar{p})$. Let $q^0 \leq \underline{p} < q^1, \ldots, \hat{p}, \ldots, q^N < \bar{p} < q^{N+1}$ denote this grid. Along the candidate equilibrium, the belief p_i evolves on this grid until it hits either q^0 or q^{N+1}, at which point no one investigates.

Let $\tilde{i}$ denote the last investigator who works: we have $p_{\tilde{i}} \in \{q^1, q^N\}$ and $p_{\tilde{i}+1} \in \{q^0, q^{N+1}\}$. Also let $\tilde{p} = p_{\tilde{i}+1}$ denote the value of the belief when learning stops under the candidate equilibrium.

To implement our candidate equilibrium consider a reward scheme where an investigator's compensation is entirely determined by his signal and the final posterior $\tilde{p}$.

For any i such that $p_i = q^k \in (\underline{p}, \bar{p})$, if i reports "H", reward him by R_H^k if $\tilde{p} = q^{N+1}$ and punish him by P_L^k if $\tilde{p} = q^0$. If i reports "L", reward him by R_L^k if $\tilde{p} = q^0$ and punish him by P_L^k if $\tilde{p} = q^{N+1}$. We will see that when the maximal rewards and punishment, R and P, are high enough, one may always choose values $\{R_\theta^k, P_\theta^k\}_{\theta \in \{``L'', ``H''\}, k \in \{1,\ldots,N\}}$ to incentivize i.

For p, q on the grid, let $\pi(p, q)$ denote the probability, given subjective belief p, that $\tilde{p} = q^{N+1}$ if the public belief variable starts at q. That is $\pi(p, q)$ is the probability that an individual with prior p assigns to p_i converging to q^{N+1} in equilibrium given that the public belief, which serves as the state variable for the equilibrium, starts at q.

If i fabricates report "H" starting from prior $p_i = q^k$, he assigns a probability $\pi(q^k, q^{k+1}$ to the public belief converging to q^{N+1}, whereas if i works and receives report "H''", his belief about the continuation equilibrium is $\pi(q^{k+1}, q^{k+1})$. Similarly, if i fabricates "L''", his belief is $\pi(q^k, q^{k-1})$, while if he works and reports "L''", his belief is $\pi(q^{k-1}, q^{k-1})$. It is straightforward to verify the inequalities

$$\pi(q^{k+1}, q^{k+1}) > \pi(q^k, q^{k+1}) \tag{1}$$

and

$$\pi(q^{k-1}, q^{k-1}) < \pi(q^k, q^{k-1}), \tag{2}$$

for all $k \in [2, N-1]$. The strictness of the inequalities comes from the fact that conditional on the true state ω, the dynamic of $\{p_j\}_{j \geq i+1}$ starting any given value of p_{i+1} is strictly increasing in ω in FOSD, as is easily checked. Therefore, the probability of hitting q^{N+1} before q^0 is strictly increasing in the belief p_i that the state is high.

For $k = 1$, the investigation stops if i reports down, so (2) only holds as an equality because the investigation stops if i reports "L", but (1) is still strict, because this report triggers further investigation. The reverse is true for $k = N$: (1) only holds as an equality while (2) is strict.

Investigator i's incentives may be computed as follows: If i fabricates evidence, the best he can get is

$$\max\{\pi(q^k,q^{k+1})R_H^k+(1-\pi(q^k,q^{k+1}))P_H^k;\pi(q^k,q^{k-1})P_L^k+(1-\pi(q^k,q^{k-1}))R_L^k\}. \tag{3}$$

The left term is i's expected payoff if he fabricates "H", and the right terms is the payoff if he fabricates "L". Since i can fabricate either evidence at no cost, his best payoff from fabrication is the maximum of these two terms.

If i works, he gets

$$z^k[\pi(q^{k+1},q^{k+1})R_H^k+(1-\pi(q^k,q^{k+1}))P_H^k]+(1-z^k)[\pi(q^{k-1},q^{k-1})P_L^k+(1-\pi(q^{k-1},^{k-1}))R_L^k] \tag{4}$$

where z^k is the probability of receiving signal "H" given belief q^k, and is equal to $z^k=Pr(``H''|q^k)=q^k\pi+(1-q^k)\times(1-\pi)$.

Therefore, i will be incentivized to work if the expression in (4) exceeds the expression in (3) by at least c.

This condition is straightforward to guarantee: choose P_H^k and P_L^k equal to some arbitrarily number, $\bar{P}$, and let $R_H^k=-P_H^k\frac{\pi(q^k,q^{k+1})}{1-\pi(q^k,q^{k+1})}$ and $R_L^k=-P_L^k\frac{1-\pi(q^k,q^{k-1})}{\pi(q^k,q^{k-1})}$. This guarantees that i's expected payoff from fabrication is zero, regardless of the outcome. From (1) and (2), his payoff from working truth is of order $\bar{P}$ and thus exceeds c, for $\bar{P}$ high enough.[4] This is true even if $k=1$ or N because, in that case there is one signal that i can send after working which yields a payoff of order $\bar{P}$, while the other signal yields 0. And the signal associated with a positive payoff arises with a probability that is bounded away from 0, since p_i lies in $(\underline{p},\bar{p})$ which is contained in the interior of $[0,1]$.

Moreover this scheme is feasible as long as the maximal reward R and and punishment P respectively exceed $\sup\{R_\theta^k:\theta\in\{L,H\},k\in\{1,\ldots,N\}$ and λ.

4 Main Result

For any $E'\in\mathcal{E}$, let $g(E')$ denote the probability that the real collection of evidence, E, strictly contains E' given that it contains at least E'.

Assumption 1 (Cheap fabrication) *The cost of fabricating any evidence m, $d(m)$ is bounded above by $d<c$.*

[4]Precisely, one can let $\bar{\pi}$ denote a strictly positive lower bound on the slack in inequalities (1) and (2) over all k's, whenever they hold strictly. Then the gain from working is of order $\bar{P}\bar{\pi}$.

Assumption 2 (Asymptotic scarcity)

$$\lim_{k \to \infty} \sup_{E':|E'|=k} g(E') = 0.$$

Assumption 3 (Bounded compensation) *i's compensation C_i is measurable with respect to $m = (m_1, \ldots)$ and takes values in a compact set $[-P, R]$.*

Assumption 4 (Discovery rate monotonicity) *If $0 < |E''| \leq |\hat{E}''|$, then $\lambda \leq \lambda_{E''} \leq \lambda_{\hat{E}''}$, where $\lambda \in (0, 1)$.*

Assumption 2 is an *asymptotic* hazard rate condition. Intuitively, it says that the probability that there is more evidence out to be found given that already k pieces have been discovered goes to zero as k becomes large. It does not rule out the possibility that some evidence discovery suggests that more evidence may be found. It requires that when an arbitrarily large number of pieces have been discovered, the probability that there exist even more pieces of evidence becomes arbitrarily small.

Assumption 4 means that the more evidence there is left to uncover, and the more likely it is that a serious investigation will uncover at least some of the uncovered evidence.

Theorem 1 *Under Assumptions 1–4, In any PBE, the report sequence m is independent of ω.*

5 Proof

Consider any informative equilibrium, i.e., one for which some agent works with positive probability, and let γ_i denote the probability that i works, given past evidence $(m_1, \ldots, m_{i-1})$. If, following some history, $\gamma_i = 0$, it is common knowledge that i's evidence is uninformative. Therefore, i's report can be removed from other agents' compensation functions without affecting their incentives. Without loss of generality, we focus on equilibria such that for each on path history, there exists some $\tilde{i} \in \mathbb{N}$ such that $\gamma_i > 0$ for all $i \leq \hat{i}$ and $\gamma_i = 0$ for all $i > \tilde{i}$. Similarly, we can rule out cases in which i fabricates evidence that cannot be produced by working: such evidence is uninformative and the equilibrium can be easily modified to get rid of this case.

The analysis focuses on an equilibrium with the above properties and will derive a contradiction in the form of a violation of the IC constraint of some agent.

Fix some ε arbitrarily small. From Assumption 2, there exists some integer $K > 0$ such that i) the ex ante probability that at least K pieces of evidence have been generated is strictly positive, ii) for any $k \geq K$, the probability that E contains strictly more than k pieces of evidence given that it contains at least k pieces is less than ε.

Let A denote the event that at most K pieces of evidence were generated. Note that $Pr(A) \geq 1-\varepsilon$. In what follows, we will often consider players' beliefs conditional on A. In particular, for any integers k and i, let f_i^k denote the probability that there remain k pieces of evidence to discover given i's information and conditional on A. By assumption, $f_0^K > 0$ and $f_i^k = 0$ for all $k > K$ and $i \in \mathbb{N}$.

We will use the following observation, whose proof is straightforward and omitted.

LEMMA 1 $\gamma_i > 0$ *only if* $Pr_i[\gamma_{i+1} > 0 | i \text{ } works] = 1$.

The intuition is clear if there is no message that i can produce after working that prompts $i+1$ (and hence, any $j > i$) to work, then i's compensation is independent of ω, conditionally on $(m_1, \ldots, m_i)$, and i's optimal strategy is to either shirk or fabricate a message m_i that maximizes his compensation.

To prove theorem 1, we will use an induction argument based on the number of remaining pieces of evidence. Since we did not impose any a priori bound on that number, we will apply the argument conditional on A, in which cases there are most K pieces, and then treat the general case.

The first result is to show that conditional on A, the probability that K pieces of evidence (i.e., the maximal number, conditional on A) remain to be found is non-increasing in equilibrium.

LEMMA 2 f_i^K *is nonincreasing in* i *along all on-path histories.*

Proof. Consider any i and report m_i. Let α_i and σ_i denote i's probability of fabricating evidence and of shirking, respectively,, conditional on $(m_1, \ldots, m_{i-1})$. Note that $\alpha_i + \gamma_i + \sigma_i = 1$, and that these three probabilities are completely pinned down by m_1^{i-1}: they do not directly depend on A.

i) Suppose first that i claims to have found some evidence, so $m_i \neq \emptyset$. By Bayesian updating, we have for $m_i \neq \emptyset$:

$$f_{i+1}^K(m_i) = \frac{f_i^K \alpha_i Pr(m_i|\alpha_i)}{\alpha_i Pr(m_i|fabricate) + \gamma_i Pr_i(m_i|work, A)}$$

That is: conditional on A, the probability that K pieces remain given i produced m_i is the probability that K pieces remained before when i inherited the case and that i fabricated the evidence m_i, divided by the probability of producing evidence m_i, which depends on whether i works or fabricates, and on the probability that m_i is produced in each case.

Rearranging the the previous equation yields

$$f_i^K - f_{i+1}^K = \gamma_i f_i^K \frac{Pr_i(m_i|work, A)}{\alpha_i Pr(m_i|fabricate) + \gamma_i Pr_i(m_i|work, A)} = \gamma_i f_i^K \frac{1}{\alpha_i \gamma_i \ell_i(m_i)} \tag{5}$$

where

$$\ell_i(m_i) = \frac{Pr(m_i|fabricate)}{Pr_i(m_i|work, A)}$$

is the likelihood that i fabricates m_i over i producing m_i for real.

ii) Suppose now that $m_i = \emptyset$, which is consistent with three possibilities: i shirked, i worked but there was no evidence left to uncover, sor i worked and there is some evidence to uncover but i was unlucky.

For each $k > 1$, let λ_i^k denote the probability that i uncovers real evidence conditional on A and on there being k pieces left to uncover. By assumption, $\lambda_i^k \geq \lambda$ and λ_i^k is nondecreasing in k. Let $\lambda_i = \frac{1}{1-f_i^0} \sum_{k=1}^K f_i^k \lambda_i^k$ denote the probability of uncovering some real evidence conditional on A and on the existence of some remaining evidence.

From Bayes' rule, we have

$$f_{i+1}^K(m_i = \emptyset) = \frac{f_i^K(\sigma_i + \gamma_i(1 - \lambda_i^K)}{\sigma_i + \gamma_i[f_i^0 + (1 - f_i^0)(1 - \lambda_i)]}$$

Rearranging yields

$$f_i^K - f_{i+1}^K = \gamma_i f_i^K \frac{\lambda_i f_i^0 + (\lambda_i^K - \lambda_i)}{\sigma_i + \gamma_i[f_i^0 + (1 - f_i^0)(1 - \lambda_i)]}. \tag{6}$$

Since $\lambda_i^K \geq \lambda_i$ and the denominator of the last equation is less than 1, we conclude that

$$f_{i+1}^K(\emptyset) - f_i^K \leq -\gamma_i f_i^K \lambda f_i^0. \tag{7}$$

Let $\mathcal{M}$ denote the set of all on-path histories truncated at any i and

$$f^K = \inf\{f_i^K : i \in \mathbb{N}, (m_1, \ldots, m_{i-1}) \in \mathcal{M}, \gamma_i > 0\}$$

Lemma 3 *There is a constant $\hat{\gamma} > 0$ such that for all $\epsilon > 0$ small enough and on-path history such that i works and $f_i^K \leq f^K + \epsilon$, $\gamma_j \leq \hat{\gamma} \frac{\epsilon}{f_i^K}$ for all $j \geq i$.*

Proof. From Lemma 2, it suffices to show that result for i. It will then follow for all $j \geq i$: either j doesn't work at all, or j works and $f_j^K \leq f^K + \epsilon$, so the proof for i will apply to j.

Focusing on i, hence, let M_i denote the set of all messages that i can produced on path, including the empty message, and for each $m_i \in M_i$ let $f(m_i) = f_{i+1}^K$ given message m_i.

There must be some messages $m_i \in M_i$ for which $f(m_i) \geq f^K$. Otherwise, by definition of f^K, $i+1$ (and hence, all $j \geq i+1$) does not work regardless of the message, and there is no incentive for i to work (cf. Lemma 1).

First, suppose that the only message for which $f(m_i) \geq f^K$ is the empty message: it means that all nonempty messages stop the investigation, i.e., no $j \geq i+1$ ever works following $m_i \neq \emptyset$. In this case, (7) implies that the drop in f^K is of order

$$\gamma_i f_i^0 f_i^K \lambda$$

We will show that f_i^0 is uniformly bounded below away from zero. Indeed, suppose that f_i^0 is less than some $\eta > 0$. Then it means that the probability of finding something is arbitrarily close to 1. So if i works and finds nothing, the probability that he was just unlucky is arbitrarily close to 1. This implies that i's expected payoff conditional on this outcome is almost identical to the expected payoff obtained if i shirked, which contradicts the optimality of working: by fabricating or shirking, i can replicate at lower cost the expected payoff following any evidence obtained from working.

Conditional on the empty message being the only message leading $i+1$ to work, f_i^0 must be bounded below by some constant $\hat{f}$ which incentivizes i to work. The decrease in f_i^K is therefore at least $\gamma_i f_i^K f_\emptyset$, where $f_\emptyset = \lambda \hat{f} > 0$. Since $i+1$ works following $m_i = \emptyset$, we have $f_{i+1}^K \geq f^K$ and, hence, $f_i^K - f_{i+1}^K \leq \epsilon$. This implies that

$$\gamma_i \leq \frac{\epsilon}{f_i^K f_\emptyset}.$$

Second, suppose that some nonempty messages prompt $i+1$ to work, and let $\underline{\ell}$ denote the infimum value of $\ell(m_i)$ across all nonempty messages prompting $i+1$ to work.

If $\underline{\ell} \geq L$ for some large constant L, any message that i sent causing $i+1$ to work is L times more likely to have been fabricated than having come from real evidence. Let b_i denote the probability that i produces a message that leads $i+1$ to work, if i works. Notice that b_i is bounded above by $1/\underline{\ell}$, because

$$b_i = \int_{M_i(i+1 \quad works)} dPr(m_i|work) = \int_{M_i(i+1 \quad works)} \frac{dPr(m_i|frabricates)}{\ell(m_i)} \leq \frac{1}{\underline{\ell}}.$$

But if b_i is negligible compared to $(c-d)/(R+P)$, then the event that i produces nonempty messages prompting $i+1$ to work is negligible from i's perspective, and thus not enough to incentivize i to work. Therefore, the only possibility is that $m_i = \emptyset$ also prompts $i+1$ to work and we are back to the previous case.

Therefore, there is a threshold ℓ^* that only depends on c, d and $(R+P)$, such that if $\underline{\ell} \geq \ell^*$ we are essentially in the first case.

Finally, consider the case $\underline{\ell} < \ell^*$. Then from (5), the decrease in f_i^K is at least $\gamma_i f_i^K \hat{f}$, where $\hat{f} = 1/(1+\ell^*)$. Since the decrease is no larger than ϵ, this implies that

$$\gamma_i \leq \frac{\epsilon(1+\ell^*)}{f_i^K}.$$

Letting $\hat{\gamma} = \max\{1+\ell^*, 1/f_\emptyset\}$ proves the lemma.

Lemma 4 $f^K = 0$.

Proof. Suppose on the contrary that $f^K \geq \delta > 0$, and consider any history where i works and $f_i^K \leq f^K + \epsilon$ for ϵ suitably small to be chosen later. From Lemma 3, $\gamma_j \leq \epsilon\frac{\hat{\gamma}}{\delta}$ for all $j \geq i$. In particular, from i's perspective, the probability that some investigation takes place at any *given* round $j > i$ is of order ϵ. Because the number of investigation rounds after round i is potentially unbounded, however, the probability of investigations in at least one round $j > i$ could a priori be non-negligible.

Let β_j denote the probability that j finds new evidence, conditional on j's information. We will prove that $E_i[\sum_{j>i} \beta_j]$ is of order ϵ. For ϵ small enough, it will show that i cannot be incentivized to work.

To prove this result, consider any $j > i$. If $\gamma_j = 0$, then clearly $\beta_j = 0$, so suppose instead that $\gamma_j > 0$ and, for any nonempty evidence $m \in M_j$ that j may produce in equilibrium, let $\gamma_j(m)$ denote the probability that j produces m conditional on A, and let $\alpha_j(m)$ denote the probability that j fabricates m (this is unconditional on A, since this only depends on the j's randomization strategy, which only depends on past reports m_1^{j-1}.[5]

We can rewrite the Bayesian updating equation (5) as

$$f_j^K - f_{j+1}^K(m) = f_j^K \frac{\gamma_j(m)}{\gamma_j(m) + \alpha_j(m)}$$

The probability that j finds evidence is

$$\beta_j = \sum_{m \in M_j \setminus \{\emptyset\}} \gamma_j(m).$$

[5]The argument is stated here when the set of pieces of evidence is countable. The argument is easy to modify by using distributions if M_j is a more general space.

Combining the previous two equations and using that $f_j^K \geq \delta$, we have

$$\beta_j \leq \frac{1}{\delta} \sum_{m \in M_j \setminus \{\emptyset\}} (f_j^K - f_{j+1}^K(m))(\gamma_j(m) + \alpha_j(m)) = \frac{1}{\delta} E_j[(f_j^K - f_{j+1}^K) 1_{m_j \neq \emptyset}] \leq \frac{1}{\delta} E_j[f_j^K - f_{j+1}^K]$$

Therefore,

$$E_i[\beta_j] \leq E_i[\frac{1}{\delta} E_j[f_j^K - f_{j+1}^K]] = E_i[\frac{1}{\delta}[f_j^K - f_{j+1}^K],$$

where the inequality follows from the law of iterated expectations.

Recalling that $\tilde{i} \geq i$ denotes the last j who works with positive probability, this person is random from i's perspective, but we have

$$E_i \left[\sum_{j>i} \beta_j \right] = E_i \left[\sum_{j=i+1}^{\tilde{i}} \beta_j \right] = E_i \left[\sum_{j=i+1}^{\tilde{i}-1} \beta_j + \beta_{\tilde{i}} \right].$$

From (3), we have $\beta_{\tilde{i}} \leq \gamma_{\tilde{i}} \leq \epsilon \frac{\hat{\gamma}}{\delta}$.

For $j < \tilde{i}$,

$$E_i \left[\sum_{j=i+1}^{\tilde{i}-1} \beta_j \right] \leq \frac{1}{\delta} E_i \left[\sum_{j=i+1}^{\tilde{i}-1} f_j^K - f_{j+1}^K \right] = \frac{1}{\delta} E_i \left[f_{i+1}^K - f_{\tilde{i}}^K \right] \leq \frac{\epsilon}{\delta}$$

Combining this yields

$$E_i[\sum_{j>i} \beta_j] \leq \epsilon \frac{\hat{\gamma} + 1}{\delta}$$

which shows that i's incentives to work are of order ϵ, contradiction the assumption that i worked, for ϵ small enough.

This proves that for all $\epsilon > 0$ small enough , there is a history and an i such that i) $f_i^K \leq \epsilon$ and ii) i works. The rest of the proof focuses on such histories where ϵ wi;ll be chosen small enough below.

We wish to prove a result similar to 4 for all $k \leq K$, and will proceed by induction.

Suppose that $K > 1$ and let f^{K-1} denote the infimum of f_i^{K-1} over all i's and histories for which $f_i^K \leq \epsilon$ and i works.

Lemma 5 $f^{K-1} = 0$.

Proof. Suppose on the contrary that $f^{K-1} > \delta$. By Bayesian updating, we have for $m_i \neq \emptyset$

$$f_{i+1}^{K-1}(m_i) = \frac{f_i^{K-1} \alpha_i Pr(m_i | fabricate) + 1_{|m_i|=1} f_i^K \gamma_i Pr_i(m_i | work, A)}{\alpha_i Pr(m_i | fabricate) + \gamma_i Pr_i(m_i | work, A)}$$

Therefore,

$$f_i^{K-1} - f_{i+1}^{K-1}(m_i) = \frac{(f_i^{K-1} - f_i^K 1_{|m_i|=1})\gamma_i Pr(m_i|work, A)}{\alpha_i Pr(m_i|fab) + \gamma_i Pr(m_i|work, A)}$$

For ϵ less than δ, the numerator is always positive, so f_i^{K-1} decreases.

For $m_i = \emptyset$, we have as before, still letting σ_i denote the probability that i shirks and $\lambda_i \geq \lambda > 0$ denote the probability conditional on A that i finds any remaining evidence if some exists if i works. From Bayes' rule, we have

$$f_{i+1}^{K-1}(m_i = \emptyset) = \frac{f_i^{K-1}(\sigma_i + \gamma_i(1 - \lambda_i^{K-1}))}{\sigma_i + \gamma_i[f_i^0 + (1 - f_i^0)(1 - \lambda_i)]}$$

Rearranging yields

$$f_i^{K-1} - f_{i+1}^{K-1} = \gamma_i f_i^{K-1} \frac{\lambda_i f_i^0 + (\lambda_i^{K-1} - \lambda_i)}{\sigma_i + \gamma_i[f_i^0 + (1 - f_i^0)(1 - \lambda_i)]}. \tag{8}$$

Noticing that $\lambda_i^{K-1} \geq \lambda_i - \epsilon \geq \lambda - \epsilon$ because

$$\lambda_i = \frac{\sum_{k=1}^K f_i^k \lambda_i^k}{\sum_{k=1}^K f_i^k}$$

and $f_i^K \leq \epsilon$ and $\lambda_i^{K-1} \geq \lambda_i^k$ for $k \leq K-1$. and the denominator is less than 1, we conclude that the drop in f^K is bounded below

$$f_{i+1}^K(\emptyset) - f_i^K \leq -\gamma_i f_i^K \lambda f_i^0. \tag{9}$$

This proves a version of Lemma 2 for $K-1$, which holds for all histories such that $f_i^K \leq \epsilon$ where ϵ is chosen to be small relative to λ. It follows that a version of Lemma 3 and, subsequently 4 holds for all such histories.

Proceeding by induction, there exists a history for which i works, but $f_i^k = O(\epsilon)$ for all integers $k \in [1, K]$. Now consider i's situation: over A, the probability that more evidence can be discovered is of order ϵ and hence negligible. But over A^c, there are two cases: either more than K pieces have already been discovered with probability arbitrarily close to 1, so the probability of discovering even more is less than ε (by our choice of A), or there is a nonnegligible probability that strictly fewer than K pieces have been discovered already, but in this case this should be the case as well if exactly K pieces were generated, by Assumption 4, which would imply that conditional on A, f_i^0 bounded above away from 1 by a fixed constant, contradicting the fact $1 - f_i^0 = \sum_{k=1}^K f_i^K$ is of order ϵ.

This proves that i has no incentive to work, which brings the desired contradiction.

6 Extensions

Joint investigation and fabrication

As the proof makes clear, there is no change in the result if investigators are allowed to fabricate/falsify/misreport evidence after a serious investigation. This result should be clear for two reasons. First, allowing this possibility only increases the set of possible deviations relative to work. Second, as the proof illustrates, what matters for the argument is whether i ever benefits from learning something about the truth, and the answer is that he does not be benefit from it. This result does not change if i can manipulate his report afterward.

Altruism towards other investigators

The result holds without change if an investigator cares about other investigators' utility, as long as the resulting utility is still uniformly bounded. Formally, such altruism would only alter the compensation function.

Intrinsic motivation

The result is also unchanged if an investigator cares about other investigators' reports, above and beyond any material reward or punishment that comes with it. For example, an investigator may be pleased if subsequent investigators follow his opinion or may take pleasure in contradicting a past investigator, making an "original" report, etc. All these considerations are already captured by a compensation function that depends on all reports

Random ordering of investigators

Suppose that the ordering of investigators is only revealed sequentially: an investigator doesn't know the identity and compensation scheme of future investigators. These considerations do not affect the analysis, because what matters for the argument is how to interpret past evidence: as long as an investigator knows the equilibrium strategy of past investigators, which is conditional on their knowledge at the time, something which the later investigator can reconstruct, he can interpret past evidence correctly.

The forward looking part of the analysis concerns only the probability that future investigators will uncover evidence, and this bound does not depend on the order of future investigators, or on their particular compensation scheme.

7 Discussion: Overturning the result

7.1 The role of ethical behavior

One of the key assumption is that an investigator's compensation does not directly depend on the true state of the world ω.

In some context, the investigator may have no material benefit from the truth, but may still care about it for other reasons. For example, a medical researcher may genuinely wish to learn more about a disease and its treatment; a police officer may genuinely care about find the real culprit of a crime, etc.

Such ethical behavior can circumvent the conclusion of Theorem 1. Moreover, these effects can be combined: for example, if investigator 2 is ethical, then not only will he contribute to learning about the fact, but his report may also be used to incentivize investigator 1 to work. The strength of these effects depends on i) the number of ethical investigators and ii) other investigators' belief about this number.

If investigators are sufficiently patient and/or punishment and rewards are sufficiently extreme, then even a small chance that some future investigators be ethical can suffice to incentivize amoral ones.

In general however, if punishment and rewards are limited, the expected number of ethical investigators must be sufficiently large to incentivize amoral investigators.

This observation leads to several issues.

First, the extent to which people behave correctly depends critically on how ethical they think others are: if no one believes in the existence of ethical agents, then it may be impossible to sustain any equilibrium in which societies and organization can learn about the type of facts studied here. In particular, public declaration undermining the public's trust in officials' standards of ethic may be particularly damaging, as may be educations purely based on the belief that agents are completely amoral.

Secondly, to the extent that ethical agents do exist, it may be critical to select them for the type of investigations studied in this paper. Law enforcement officers, prosecutors, judges, scientists, and experts of various sorts may need to exhibit an ethical sense. It may be important to divert amoral individuals away from these professions and from this perspective, lucrative professions which appeal to greedy or amoral individuals may paradoxically serve a useful purpose, by diverting these individuals from the type of tasks where moral behavior is needed.

Material reasons to care about the truth

When interactions are repeated, investigators may care about the truth per se even if they are self interested. For example, someone asked to learn about a disease may care that this disease affect him in the future, which makes the dependence of the truth explicit in his objective function. Similarly, an executive trying to learn about the strengths and weaknesses of his firm may value the answer because of its direct relevance for the firm profitability in the future

Such situations stands in sharp contrasts with a criminal investigation where the police officer assigned to the case has presumably little direct material interest in learning who committed the crime, or a judicial opinion where the judge is very unlikely to be materially affected by the verdict on a defendant.

8 Literature review

TBA

Inspection games

Costly monitoring

Monitoring the monitor

Herding models

Law enforcement

Folk theorem in dynamic games

Ethical behavior

Parallel monitoring