Costless Information and Costly Verification: A Case for Transparency

Deniz Kattwinkel

University College London

Jan Knoepfle

Queen Mary University of London

A principal has to take a binary decision. She relies on information privately held by an agent who prefers the same action regardless of his type. The principal cannot incentivize with transfers but can learn the agent's type at a cost. Additionally, the principal privately observes a signal correlated with the agent's type. Transparent mechanisms are optimal: the principal's payoff is the same as if her signal was public. A simple cutoff form is optimal: favorable signals ensure the agent's preferred action. Signals below this cutoff lead to the nonpreferred action unless the agent appeals. An appeal always triggers type verification.

I. Introduction

A principal has to take a binary decision. Her preferences are determined by an agent's private type. The agent prefers one of the two actions independent of his type. The principal cannot use monetary transfers to

We are grateful to Marina Halac, Stephan Lauermann, Benny Moldovanu, and Sven Rady for continual guidance and support. We likewise thank Ian Ball, Martin Cripps, Martin Hagen, Daniel Hauser, Philippe Jehiel, Andreas Kleiner, Daniel Krähmer, Konrad Mierendorff, Pauli Murto, Martin Pollrich, Ludvig Sinander, Andre Speit, Dezsö Szalay, Tymon Tatur, Juuso Välimäki, and seminar audiences at the 11th BiGSEM (Bielefeld Graduate School of Economics and Management) Workshop on Economic Theory, the Jerusalem Summer

Electronically published January 20, 2023

Journal of Political Economy, volume 131, number 2, February 2023.

https://doi.org/10.1086/721618

incentivize the agent to reveal his type. Instead, she privately observes a costless imperfect signal about the agent's type and can verify his exact type at a cost.

Applications covered by this setting include (1) a human resources department decides whether to hire a candidate, (2) a judge decides whether to acquit or convict a defendant, and (3) a competition authority decides whether to grant or deny a company permission to acquire or merge with another firm.

The agent has a clear preference toward one action: the candidate wants to be hired, the defendant wants to be acquitted, and the company wants to merge. The principal's preferences depend on the agent's private type: the candidate's ability, the defendant's guilt, or the company's competitive position in the market.

In many environments, monetary transfers to elicit the agent's type are not feasible;[1] beyond the above-mentioned applications, monetary incentive schemes are rarely observed in bureaucracies and public sector decisions. In these environments, the principal can often learn the agent's type through verification, for example, conducting an assessment center, a trial, or a market analysis. Verification is costly, so the principal wants to economize on it.

Typically, costly information acquisition is not the only way to learn about the agent's private type. The human resources department receives references from previous supervisors, the judge sees the outcome of pretrial investigations, and the competition authority has sector-specific knowledge from its supervisory function. That is, the principal privately observes a signal about the agent's type.

Private information and costly verification are natural instruments in bureaucracies when monetary transfers are absent. In this paper, we want to explore how the principal can use her private information to minimize verification costs and decision inefficiencies. To this end, we analyze principal-optimal mechanisms in this setting.

As the principal's signal is private, the agent forms a belief over the signal realization. As signal and type are correlated, this belief varies with the agent's type. Can the principal benefit from secrecy to screen the agent

School in Economic Theory 2017, the Stony Brook Conference on Game Theory 2018, the ESEM (Econometric Society European Meeting) 2018 (Cologne), SAET (Society for the Advancement of Economic Theory) 2021, Games 2020.1 (Budapest), and at Aalto, Bonn, Yale, and Queen Mary. Both authors thank the Yale Economics Department for their hospitality during the respective research visits, when parts of this paper were written. This work was supported by the German Research Foundation (DFG) through CRC TR 224 (Collaborative Research Center Transregio 224; project B04). We thank the editor and anonymous referees for insightful suggestions that significantly improved the paper. This paper was edited by Emir Kamenica.

[1] The assumption is that payments cannot depend on the agent's report. Even though a public sector job entails payments, if the payment is fixed, it cannot be used to incentivize truthful reports of the candidate's ability.

along the various beliefs, or should she make her private signal public before interacting with the agent? This question is relevant in practice, as it determines whether secret procedures perform better than transparent procedures. In sharp contrast to settings in which money can be used to screen the agent, we show that transparent procedures are optimal.

Results.—The principal's benefit from choosing the agent-preferred action over the alternative increases in the type. We show that among all Bayesian incentive-compatible (BIC) mechanisms, it is optimal for the principal to commit to a simple *cutoff-with-appeal* procedure: if the signal makes her sufficiently certain that the agent's type is high, she takes his preferred action without communicating with him. If the signal falls below this cutoff, she takes the nonpreferred action by default but gives the agent the possibility to appeal. Upon appeal, the principal verifies the agent's type and revises her decision whenever the type exceeds a certain threshold. The optimal appeal threshold is such that for all types above the threshold, the principal's benefit from revising her decision exceeds the verification cost.

An important feature of this mechanism is that it does not require the principal's signal to be secret. Therefore, the principal cannot profit from strategically hiding or releasing parts of her information. This makes the case for transparency. For the applications mentioned above, this implies that the human resources department showing the references to the candidate, the judge informing the defendant of pretrial investigation results, or the competition authority publicizing her market assessments does not constrain the implementation of optimal procedures.

The question whether procedures and institutions should be transparent has gained a lot of attention in the past decades in politics, corporate governance, and law; see Bushman, Piotroski, and Smith (2004), Prat (2005), and references therein for contributions by economists. Transparent procedures have many advantages, most prominently accountability and predictability. Considering these advantages, any nontransparent procedure must be justified. Often, these justifications invoke efficiency concerns, as in Prat (2005). Our paper contributes to this discussion by showing, for a natural class of problems, that transparency does not conflict with efficiency. Therefore, efficiency cannot justify secrecy.[2] Our findings are in line with the continual advancement of transparency in the private and public sectors. One example is the evolution of codes of criminal procedure in continental Europe. While modern codes prescribe the disclosure of all potential charges to the defendant, this was not always the case.[3]

Moreover, the cutoff-with-appeal mechanism is deterministic. Together with transparency, this absence of randomization implies that the procedure

[2] One alternative justification beyond our model could be privacy concerns. For a broad discussion, see Hood and Heald (2006).

[3] See sec. III.B for a discussion of the development of the Austrian code of procedure.

is fully predictable for the agent. Once he learns the signal, he can perfectly predict the outcome. This enables the public to hold institutions accountable in case of deviations.

As the first step in the analysis, we combine a revelation principle with optimality arguments to show that it is optimal for the principal to use a truthful direct mechanism. A direct mechanism specifies, for any combination of type report and signal realization, a probability of verification and a probability that the agent-preferred action is taken. Whenever verification reveals that the agent's reported type is different from his true type, he is punished with the nonpreferred action. In the remainder, we take the principal's decision to be an allocation choice (where the agent-preferred option is to allocate). Since the agent does not know the signal realization, from his perspective, every report leads to an allocation lottery. A direct mechanism is BIC if reporting his true type maximizes the agent's expected payoff.

The standard approach under independence—to characterize mechanisms in their reduced form with interim expectations—does not apply to our setting because different types hold different beliefs about the signal realization. This belief heterogeneity requires new techniques to characterize optimal mechanisms. We show that it is without loss for the principal to use transparent procedures. In a transparent procedure, publishing the signal eliminates the belief heterogeneity.

The optimality of transparency stands in stark contrast to existing results on mechanism design with transfers where belief heterogeneity is used to eliminate agency costs.[4] To explain this contrast, we illustrate how the principal could potentially exploit secrecy to increase efficiency. We present mechanisms that can be implemented under secrecy but are infeasible when the signal is public. However, our main conceptual contribution shows that the *optimal* mechanism does *not* make use of secrecy. It satisfies three properties that eliminate the benefit from secrecy: it is deterministic, and the allocation is pointwise increasing in the agent's type and pointwise increasing in the principal's signal. Our transparency result follows by verifying that these properties indeed eliminate any benefit from secrecy. To the best of our knowledge, this is the first result of this kind under correlated information.[5]

In many applications the verification technology may be imperfect. Take the hiring example mentioned above. First, a low-ability type may sometimes achieve a high test score by luck. Second, the variety of ability types of a job candidate may exceed the available degrees and certificates

[4] See Myerson (1981), Crémer and McLean (1988), and McAfee and Reny (1992). See sec. VI for a discussion of the related literature.

[5] While Crémer and McLean (1988) demand dominant incentive compatibility in the second stage, the surplus extraction in the first stage crucially exploits that agents do not know the realization of the correlated information.

that can be verified. We call the first form of imperfection *noisy* verification and the second form, in which different types remain indistinguishable in the verification process, *coarse* verification. In both cases, we characterize optimal transparent mechanisms. Under noisy verification, the optimal transparent mechanism shows small interior allocation probabilities for some types, violating the nonrandomness property. As a consequence, the principal can beneficially shift these small probabilities across signals under secrecy. This benefit from secrecy is negligible for small noise in the verification. In particular, it vanishes linearly in the verification error. Under coarse verification, the optimal transparent mechanism satisfies the three properties above—monotonicity in type and signal and nonrandomness—and is also optimal in the larger class of secret mechanisms. There is no loss from transparency.

In the online appendix, we analyze the cases when the agent also incurs a cost from verification and when the principal's signal has a direct effect on her preferences. Again, the three properties introduced above help us to determine whether transparency is optimal. In the case when the agent's verification cost makes false appeals less attractive, the principal can save costs by slightly lowering the verification probability. The resulting mechanism is not deterministic, and the principal can benefit from secrecy. This benefit is negligible for small agent costs. In particular, it vanishes linearly in the agent's cost. Finally, suppose that the principal's signal has a direct effect on her payoffs. When this effect is positive, our result carries over: transparency comes without loss for the principal. If, in contrast, the direct effect is negative, the principal benefits from secrecy because the optimal transparent mechanism is decreasing in the signal.

Section II contains the model and the revelation principle (theorem 1). Section III presents the optimal mechanism (theorem 2). The proof is contained in section IV: proposition 1 establishes our transparency result, and proposition 2 characterizes optimal transparent mechanisms. In section V, we analyze imperfect verification, and section VI discusses related literature.

II. Model

A. *Setup*

The principal (she) decides whether to allocate a single, indivisible good to the agent (he). Her allocation preferences depend on the agent's private type t. The set of possible types T is finite and ordered.

While t is unknown to the principal, she receives costless information about it in form of a private signal $s \in S$, finite and ordered. Type t and signal s are jointly distributed with probability $f(t, s) > 0$ for all $t \in T$, $s \in S$. The signal satisfies the monotone likelihood ratio property (MLRP):

for all $t < t' \in T$, the ratio $f(t', s)/f(t, s)$ is increasing in s. This implies that a higher signal is more indicative of a higher type.[6] In addition to the costless information, the principal has the option to learn t at verification cost $c > 0$. Verification is perfect; she learns the exact type.[7]

The principal derives valuation $v(t) \in \mathbb{R}$ when allocating the good to type t. The value she derives from not allocating is 0. Therefore, v represents the net value for the principal. Valuation $v(t)$ is nondecreasing in t, and there are $t', t'' \in T$, with $v(t') < 0 < v(t'')$.[8] When the agent has type t, he receives utility $u(t) > 0$ from the good. His utility from not receiving the good is always zero.

B. Mechanisms

We study the interaction between principal and agent in a mechanism-design setting and characterize mechanisms that maximize the principal's expected valuation net of verification costs. The principal can design arbitrary mechanisms, and the agent plays a Bayesian best response after learning his type. A key question in this setting is, Can the principal use her private information—the signal—to elicit the agent's information? We give the principal maximal flexibility to use her information and assume that the signal is contractible. She can commit to mechanisms that are contingent on the signal realization.[9] In the appendix, we define a broad class of dynamic mechanisms that allow the principal to release garblings of her information at any point of the interaction. This covers any potential for information design by the principal. Theorem 1 shows that this can be captured without loss within a simple class of mechanisms.

A *direct* mechanism specifies for any type-signal pair (t, s) two probabilities, $x(t, s)$ and $z(t, s)$, and proceeds as follows. It asks the agent to report his type. On the basis of this report t and the signal realization s, one of three distinct events occurs:

1. with probability $x(t, s)$, the good is allocated to the agent, and he is not verified;
2. with probability $z(t, s)$, the agent is verified; then, the good is allocated to him if and only if he is found to have reported truthfully;

[6] The MLRP is equivalent to requiring that t and s be affiliated.

[7] Whether the verification technology reveals the true type of the agent or just confirms for a specified type whether the agent has this type or not does not alter our results.

[8] Otherwise, the principal can implement the optimal allocation without the agent's information.

[9] She can commit to truthfully communicate her signal realization to the mechanism. An alternative approach would be to consider the informed-principal problem, requiring the mechanism to make truthful communication incentive compatible for her. We show in sec. VI that our results are robust to this modeling choice. The optimal mechanism constitutes an equilibrium in the informed-principal problem.

3. with probability $1 - x(t, s) - z(t, s)$, the good is not allocated to the agent, and he is not verified.

Feasibility requires the total allocation probability $x(t, s) + z(t, s)$ not to exceed 1. In the remainder, we refer to $x(t, s)$ as the *nonverified allocation* probability. A mechanism is called *truthful* if reporting truthfully is a best response for all types. The theorem below combines a revelation principle with optimality considerations.

THEOREM 1. There is a direct truthful mechanism that maximizes the principal's expected valuation net of verification costs.

In the proof in appendix A, we first derive a revelation principle—reminiscent of Ben-Porath, Dekel, and Lipman (2014) and Akbarpour and Li (2020)—for our setting with correlated information. Then, we exploit that any optimal mechanism has to satisfy two intuitive properties: (1) *maximal punishment*: after verification reveals a misreport, the agent does not receive the good; and (2) *minimal verification*: after his report is verified to be true, the agent receives the good for sure.

Truthful direct mechanisms do not restrict the principal's ability to strategically release information. Take a mechanism that is not in direct form. Suppose that the mechanism reveals a garbling of the signal and then asks the agent to send a message. Different realizations of the garbling induce different beliefs about the signal before the agent sends his message. Theorem 1 shows that this information-design mechanism can be replicated by a direct mechanism.[10] The direct mechanism asks the agent for his type and then internally simulates the original mechanism, with the agent's best response corresponding to the reported type. Although the agent does not receive the information from the garbling before his report in a direct mechanism, when evaluating the expected utility from different reports, he takes the perspective of the simulated agent. Through this channel, the information design in the original mechanism affects the incentives of the agent to report truthfully in the direct mechanism.

C. *Incentive Compatibility*

In standard mechanism-design problems, the set of feasible allocations is pinned down by the incentive-compatibility (IC) constraints through the integral characterization by Myerson (1981). Our design setting is nonstandard in two ways: the absence of transfers and the presence of correlated information. Ben-Porath, Dekel, and Lipman (2014) show that the absence of transfers impedes the integral characterization, and they present another tractable characterization for this case when information is independently distributed. In contrast to our setting, there, interim

[10] We discuss this formally in the appendix (see corollary 1).

expectations are sufficient to characterize payoffs. Correlated information impedes this approach (see example 1 below) and requires an alternative methodology.

1. Bayesian Incentive Compatibility

Absent monetary transfers, the agent cares solely about the probability of receiving the good. Consider the incentives of an agent of type t. He does not know the signal realization. If he reports truthfully, he faces the random allocation probability $x(t, s) + z(t, s)$, which depends on the random variable s. Whether his report is verified is irrelevant for him. If, however, type t makes a report $\hat{t} \neq t$, he receives the good with random probability $x(\hat{t}, s)$, that is, only if he is not verified. Therefore, type t prefers reporting t to reporting $\hat{t}$ if

$$u(t) \cdot \mathbb{E}_s[x(t, s) + z(t, s)|t] \geq u(t) \cdot \mathbb{E}_s[x(\hat{t}, s)|t].$$

Since every type derives strictly positive utility from the good $(u(t) > 0)$, type t's preference intensity can be eliminated from the IC constraint. The agent simply maximizes his expected allocation probability, and the Bayesian incentive constraints can be expressed as follows: for all $t, \hat{t} \in T$,

$$\mathbb{E}_s[x(t, s) + z(t, s) - x(\hat{t}, s)|t] \geq 0. \qquad (\text{BIC}_{t,\hat{t}})$$

In our model with correlated information, different types hold different beliefs about s, so that the interim expected allocation probability from a given report is different for different types. Hence, with correlation, one must consider the interim expectations for all combinations of true type and report. This is different with independent information, where the expected utility of any misreport can be treated as independent of the true type. By independence, in Ben-Porath, Dekel, and Lipman (2014), incentive compatibility holds whenever the type with the lowest expected allocation probability does not want to misreport. The following example illustrates that this does not hold with correlation; the example will be revisited below to build intuition for our results.

Example 1. The agent's type is either high or low, $t \in \{L, H\}$. The principal observes signal $s \in \{\ell, h\}$. Type and signal are jointly distributed according to

$$\begin{pmatrix} f(L,\ell) & f(L,h) \\ f(H,\ell) & f(H,h) \end{pmatrix} = \frac{1}{6}\begin{pmatrix} 2 & 1 \\ 1 & 2 \end{pmatrix}. \qquad (1)$$

Consider the following nonverified allocation probabilities,

$$\boldsymbol{x} = (x(L,\ell), x(L,h), x(H,\ell), x(H,h)) = \Big(\underset{(L,\ell)}{0}, \underset{(L,h)}{1}, \underset{(H,\ell)}{1}, \underset{(H,h)}{0}\Big), \qquad (2)$$

which allocate if and only if report and signal do not match. Suppose that there is no verification. This mechanism is not incentive compatible. Both types prefer to misreport because they put higher probability on the matching signal. For both types, the interim expected allocation probability from truth telling is 1/3, while misreporting yields 2/3.

A general characterization of incentive compatibility with belief heterogeneity remains an open question. Instead of characterizing all incentive-compatible mechanisms explicitly, Crémer and McLean (1988) and McAfee and Reny (1992) show that, with money, the belief heterogeneity allows the principal to extract all surplus from the agents.[11] Since full surplus extraction is not feasible in our setting, we consider the following, more restrictive incentive constraints as an intermediate step toward characterizing optimal mechanisms.

2. Ex Post Incentive Compatibility (EPIC) and Transparency

We call a mechanism *transparently* implementable if there is an implementation that starts with the principal making her information public. For direct mechanisms, transparency requires that, after observing any signal s, all types t report truthfully. This is, thus, equivalent to requiring EPIC: for all $t, \hat{t} \in T$ and for all $s \in S$,

$$x(t, s) + z(t, s) - x(\hat{t}, s) \geq 0. \qquad (\text{EPIC}(s)_{t,\hat{t}})$$

Every Bayesian incentive constraint $(\text{BIC}_{t,\hat{t}})$ is a weighted sum of the corresponding $(\text{EPIC}(s)_{t,\hat{t}})$ constraints. Therefore, every incentive-compatible transparent mechanism is also BIC. In example 1, after learning the signal, both types agree which report is most profitable (the report contrary to the signal). With transparency, when all types learn the signal, the belief heterogeneity is resolved. This facilitates the characterization of optimal transparent mechanisms (sec. IV.B). Paired with our main conceptual contribution (sec. IV.A)—that any BIC mechanism can be made transparent without loss for the principal—this yields optimal mechanisms in the larger class of BIC rules.

III. Optimal Mechanisms

The principal designs a mechanism that maximizes her expected utility from the allocation net of the cost of verification. If the good is assigned without verification, she gains $v(t)$. In the case of allocation with prior verification, she additionally pays cost c. Hence, the principal's problem can be stated as the following linear program:

[11] They establish the existence of an incentive-compatible mechanism that allocates efficiently and extracts all surplus and, therefore, must be optimal.

$$\max_{(\boldsymbol{x},\boldsymbol{z})\geq 0} \mathbb{E}[x(t,s)v(t) + z(t,s)(v(t) - c)],$$
$$\text{subject to } \forall t, \hat{t} \in T: (\text{BIC}_{t,\hat{t}}) \quad \text{(LP)}$$
$$\text{and } \forall t \in T, s \in S: x(t,s) + z(t,s) \leq 1.$$

Note that the principal optimizes subject to the Bayesian incentive constraints (not to the stronger transparency constraints). The principal's value from an incentive-compatible transparent mechanism cannot exceed the value from the above problem. The following class of mechanisms plays an important role in the ensuing analysis:

Definition 1. A mechanism $(\boldsymbol{x}, \boldsymbol{z})$ is called "cutoff with appeal" if there exists a cutoff $\bar{s}$ and an appeal threshold $\bar{t}$ such that

(i) if $s \geq \bar{s}$, then $x(t,s) = 1$ for all t and $z(t,s) = 0$ for all t;

(ii) if $s < \bar{s}$, then $x(t,s) = 0$ for all t and $z(t,s) = \begin{cases} 1 & \text{for } t \geq \bar{t}, \\ 0 & \text{for } t < \bar{t}. \end{cases}$

Figure 1 sketches a cutoff-with-appeal mechanism. If the signal realization is above the cutoff $\bar{s}$, the principal allocates the good to the agent irrespective of his reported type and without verification ($x = 1$). If the signal is below the cutoff, the agent can receive the good only after his type report is verified ($z = 1$) to be above the threshold $\bar{t}$. Note that appeal threshold $\bar{t}$ is the same after all signals. We call this class "cutoff with appeal" because it can be implemented by the following procedure. For signals above $\bar{s}$, the principal allocates without eliciting any information from the agent. For signals below $\bar{s}$, the default is not to allocate, but the principal gives the agent the opportunity to appeal. An appeal is granted only after the type is verified to be above the threshold. Our main result shows that the optimal mechanism can be found in this class and specifies the optimal cutoff for the signal and the optimal appeal threshold for the type.

Theorem 2. The principal's problem is solved by the cutoff-with-appeal mechanism with cutoff $\bar{s}$ and appeal threshold $\bar{t}$ given by[12]

$$\bar{s} = \min\{s | \mathbb{E}_t[v(t)|s] > \mathbb{E}_t[(v(t) - c)^+|s]\} \quad \text{and}$$
$$\bar{t} = \min\{t | v(t) - c > 0\}.$$

By the positive correlation (MLRP), higher signals make the principal more optimistic about the agent's type. If the signal exceeds the cutoff $\bar{s}$, she is sufficiently optimistic and allocates without eliciting any further information from the agent. If the signal is below the cutoff, the principal is pessimistic and makes the allocation type dependent. To prevent misreports,

[12] For $r \in \mathbb{R}$, we denote the positive part by $(r)^+ = \max\{r, 0\}$.

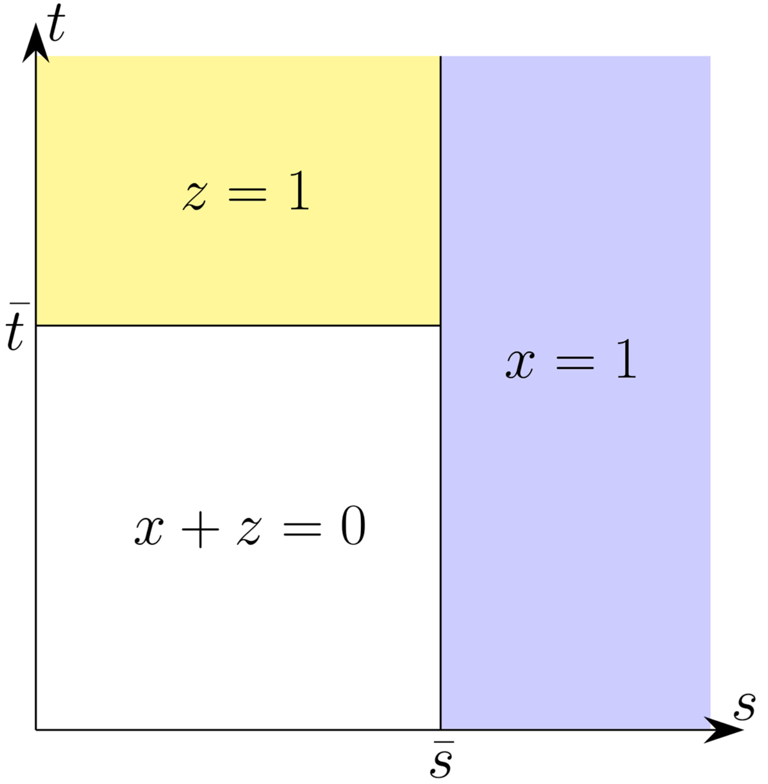

FIG. 1.—Cutoff with appeal.

she allocates only after type verification. The threshold $\bar{t}$ is such that for all higher types the principal profits from allocating, even accounting for the verification costs (i.e., when $v(t) - c > 0$). Therefore, given a signal s below the cutoff, the principal's expected value is $\mathbb{E}_t[(v(t) - c)^+ | s]$. The optimal cutoff $\bar{s}$ is set such that the principal prefers this value when the signal falls below $\bar{s}$ and prefers the expected value from allocating to all types $\mathbb{E}_t[v(t)|s]$ otherwise.

In this optimal mechanism, the principal does not exploit the heterogeneous beliefs for information elicitation. In particular, the agent would report his type truthfully, given any belief about the signal. Hence, the cutoff-with-appeal procedure does not entail complex surplus-extracting schemes as the literature suggests for optimal mechanisms in settings with correlation and money (Crémer and McLean 1988). This does not mean that benefiting from belief heterogeneity is generally impossible in our setting (see sec. IV.A). Theorem 2 shows that when the objective is to allocate efficiently, the transparent cutoff-with-appeal mechanism outperforms these complex procedures that would require secrecy.

The proof of theorem 2 is outlined in section IV and consists of two steps. The first step contains the main conceptual contribution of the paper: the principal can achieve her optimum in the class of BIC mechanisms with a transparently implementable mechanism. Thus, transparency entails no loss. The second step completes the proof by showing that the cutoff-with-appeal mechanism in theorem 2 is optimal in the class of transparent mechanisms. Before presenting the proof, we collect important features of the optimal mechanism and demonstrate how our result applies to the optimal design of court procedures.

A. Features

Transparency.—The cutoff-with-appeal mechanism can be implemented transparently. The principal could first reveal her signal to the agent and then ask him to report his type. To see why, consider figure 1. If the signal exceeds the cutoff $\bar{s}$, the allocation is independent of the report (blue area). If the signal is below $\bar{s}$, the agent can get the good only after being verified (yellow area), so that misreporting cannot be beneficial even when the agent knows the signal. In a transparent mechanism, the agent does not have to form beliefs about the principal's signal to determine his best response. Thus, the mechanism remains incentive compatible for any potentially misspecified prior.

Predictability.—The optimal cutoff-with-appeal mechanism does not require randomization. Allocation and verification probabilities take only values in $\{0, 1\}$. Therefore, when the signal is made transparent, the decision-making process is fully predictable for the agent. Hence, the public can hold the principal accountable for violations of the announced procedure. This promotes commitment power, as covert deviations from the established rules are not possible.

Communication.—In the cutoff-with-appeal mechanism, there is communication from the agent to the principal. The principal does not learn only from her signal and from verification. To contrast, consider the following no-communication mechanism, which would be optimal if the principal could not exchange messages with the agent. There are three signal regions. If the signal is high, the principal allocates without verification. If the signal is intermediate, the principal verifies the agent's type and allocates if and only if $v(t) > 0$. If the signal is low, the principal never allocates or verifies.[13] The cutoff-with-appeal mechanism improves upon this through communication in two instances. (i) When the signal is intermediate and the agent reports that his type is below the threshold,

[13] We thank a referee for suggesting this illuminating example.

so that $v(t)$ is below c, it saves verification costs. (ii) When the signal is low but the agent reports a type above the threshold, the principal's value of allocation net the verification cost is positive, so it is beneficial for the principal to verify and allocate. The communication in the optimal mechanism can be reduced to a binary message after the principal provisionally decides not to allocate: {appeal, not appeal}.

Unique implementation.—If his type is below the threshold, the agent's chances of getting the good are unaffected by his report. Hence, he is indifferent between appealing and not appealing. To make not appealing the unique best response, consider the following amendment of the optimal mechanism. The principal offers a small allocation probability $\epsilon > 0$ when the agent does not appeal and the signal falls below the cutoff. Then, the agent has strict incentives to appeal truthfully for any type. By choosing ϵ arbitrarily small, the cutoff-with-appeal mechanism achieves the supremum payoff among all mechanisms in which the agent has a unique best response.[14] The amended mechanism above is also the optimal transparent mechanism when the verification fails with small probability ϵ (see sec. V).

Futility of information design.—The contractability of the signal gives the principal maximal flexibility to use her private information. Nevertheless, the optimality of a transparent mechanism implies that the principal does not profit from persuading the agent to reveal the truth through any form of information design. With less flexibility, for example, when the signal is not contractible, information design remains futile. Hence, our mechanism solves the informed-principal problem (see sec. VI).

B. *Informing the Defendant*

Consider the following application of our model. A judge has to decide whether to acquit or convict a defendant. The defendant privately knows whether he is guilty or innocent. When the defendant is charged, the judge observes the result of a prior investigation and decides whether to conduct a full trial, which will reveal whether the defendant is guilty at a cost. The judge wants to acquit the defendant if and only if he is innocent, while the defendant prefers being acquitted irrespective of his guilt. When we model the defendant as the agent, the judge as the principal, and the decision to acquit as the allocation ($x = 1$), we can identify the agent being innocent with $t = 1$ and the agent being guilty with $t = 0$. Capture the preferences of the judge by $v(t) = t - 1/2$.[15]

[14] This set of unique-best-response mechanisms is not closed, and it has no maximizer.

[15] If an innocent agent ($t = 1$) is convicted ($x = 0$), the net utility loss is given by $0 - v(1) = -1/2$. If a guilty agent ($t = 0$) is acquitted ($x = 1$), it is given by $v(0) = -1/2$.

The optimal mechanism in theorem 2 resembles the proceeding of a pretrial. The case is dismissed if the signal for the defendant's innocence is strong enough, that is, the charge is weak. If the signal for innocence is below this cutoff, the agent can plead guilty and is convicted, or he can request a trial by pleading not guilty, after which he is acquitted if indeed found to be not guilty and convicted otherwise.

An important implication of our transparency result is that the justice system does not profit from keeping the discovery of pretrial investigations secret. Transparent discovery is established practice in modern codes of procedures, but this was not always the case. See *Brady v. Maryland* (373 U.S. 83, decided 1963) for the case of US federal law,[16] or consider today's Austrian criminal code of procedure (StPO 1975, §6 (2)) and the code of 1803 (Franz II 1803, §331). While the modern code guarantees the defendant's right to learn about all potential charges, the version from 1803 grants the court much more discretion in the extent of information released to the defendant, stating that he has to be informed only as far as necessary to notify him that he is accused.[17]

If one extends the model and allows for more types of the agent—that is, guilty ($t = 0$), guilty of a minor crime ($t = 3/4$), and innocent ($t = 1$)—another feature of court proceedings arises in our optimal mechanism. When a defendant who is guilty of the minor crime triggers a trial by misrepresenting his type as innocent, he cannot hope for acquittal in the trial. The judge is committed to punish her for having lied, even if this is ex post inefficient ($v(3/4) = 1/4 > 0$). This feature is reflected by the harsh punishments that lying in court usually entails.

IV. The Proof of Theorem 2

A. The Case for Transparency

This subsection presents the main conceptual insight of our paper: the principal cannot exploit the secrecy of a private signal that is correlated with the agent's type to reduce his information rents. This stands in marked contrast to settings with transfers, where secrecy and correlation permit full elimination of information rents.

[16] The prosecution did not inform the defendant Brady of his companion's previous confession to the actual killing. The US Supreme Court ruled that "the government's withholding of evidence that is material to the determination of either guilt or punishment of a criminal defendant violates the defendant's constitutional right to due process."

[17] This is in line with the broader development in continental Europe from medieval inquisitorial proceedings with secret charges to modern criminal law proceedings (Kittler 2003). The most famous defendant whose charges are kept secret may be Josef K., the protagonist in Franz Kafka's novel *Der Proceß* (*The Trial*). In fact, Kittler (2003) suggests that Kafka, who was a legal scholar, based his *Proceß* not on his contemporary but the medieval proceeding standards.

PROPOSITION 1. It is without loss of optimality for the principal to use a transparent procedure. Formally, consider any mechanism $(\boldsymbol{x}, \boldsymbol{z})$ that is feasible in the principal's problem. There exists a feasible mechanism $(\tilde{\boldsymbol{x}}, \tilde{\boldsymbol{z}})$ that satisfies $(\text{EPIC}(s)_{t,\hat{t}})$ for all $s, t, \hat{t}$ and delivers a payoff to the principal no lower than $(\boldsymbol{x}, \boldsymbol{z})$.

The formal proof of this proposition can be found in appendix A. There are two channels through which the principal can potentially exploit secrecy. We revisit example 1 to illustrate for each channel (i) how it allows the principal to lower the verification costs required to implement some allocation rules and (ii) why the optimal allocation rule renders this channel futile.

EXAMPLE 1 CONTINUED (*a*). Consider the environment from example 1 with distribution as in equation (1). The cost-minimal transparent verification schedule to implement the total allocation in equation (2) verifies with probability one whenever the good is allocated, that is,

$$\boldsymbol{z} = \left(\underset{(L,\ell)}{0}, \underset{(L,h)}{1}, \underset{(H,\ell)}{1}, \underset{(H,h)}{0} \right).$$

The agent's ex post incentive constraint is binding when his type and the signal match (and he does not get the good), and it is slack otherwise. Under secrecy, the total allocation in equation (2) is optimally implemented by the verification schedule

$$\boldsymbol{z} = \left(\underset{(L,\ell)}{0}, \underset{(L,h)}{0.5}, \underset{(H,\ell)}{0.5}, \underset{(H,h)}{0} \right).$$

This creates verification costs only half of those with transparency.

With the second, cheaper verification schedule, the allocation is not transparently implementable. If type L knows that the signal is ℓ, he can get allocation with probability $1/2$ by misreporting H. In this example, secrecy allows the principal to "reuse" excess allocation probability across different signals when the agent is unaware of its realization. Why does this not work in the optimal mechanism? When the total allocation probability is nondecreasing in t for all s, an improvement as above is not possible. First, monotonicity implies that only upward incentive constraints matter.[18] Second, if an $\text{EPIC}(s)_{t,\hat{t}}$ constraint at a signal s binds for some type t, under monotonicity it must also bind for all lower types $t' < t$ at this signal. It follows that, in an optimal transparent mechanism, the lowest type's EPIC constraint must bind at all signals, so no slack can be reused under

[18] Indeed, the first step of the proof consists of presenting a relaxation of the principal's problem discarding, among others, all downward incentive constraints. Establishing monotonicity directly is complicated with belief heterogeneity.

secrecy. Note that the gain in flexibility to reuse excess utility across signals did not depend on correlation. This advantage of Bayesian versus ex post implementation is present for general distributions. In the case of independent types, however, it is precisely the pointwise monotonicity of allocations in types that leads to the BIC-EPIC equivalence (Manelli and Vincent 2010; Gershkov et al. 2013; Ben-Porath, Dekel, and Lipman 2014).

With correlation, monotonicity in t is not sufficient to conclude that transparency is optimal. The belief heterogeneity creates an additional channel to benefit from secrecy. The following two examples illustrate how the principal can exploit this when the total allocation is either nonmonotone in s or not deterministic.

Example 1 continued (b). Consider the total allocation probabilities

$$\boldsymbol{x} + \boldsymbol{z} = \left(\underset{(L,\ell)}{1}, \underset{(L,h)}{0}, \underset{(H,\ell)}{1}, \underset{(H,h)}{1} \right). \tag{3}$$

The optimal verification probabilities to implement this transparently are

$$\boldsymbol{z} = \left(\underset{(L,\ell)}{0}, \underset{(L,h)}{0}, \underset{(H,\ell)}{0}, \underset{(H,h)}{1} \right).$$

Given the distribution in equation (1), this results in verification costs of $(2/6) \cdot 1 \cdot c = (1/3)c$. Without observing the signal, the agent updates his belief conditional on his type. With the distribution in equation (1), type L's subjective belief on signal ℓ is $2/3$. Consider L's Bayesian incentive constraint:

$$\frac{2}{3} \cdot 1 + \frac{1}{3} \cdot 0 \geq \frac{2}{3} \cdot (1 - z(H, \ell)) + \frac{1}{3} \cdot (1 - z(H, h)). \tag{$\mathrm{BIC}_{L,H}$}$$

The principal can exploit that type L puts more weight on signal ℓ and shift verification probability from the type-signal combination (H, h) to (H, ℓ). Under secrecy, the allocation above is optimally implemented with verification probabilities

$$\boldsymbol{z} = \left(\underset{(L,\ell)}{0}, \underset{(L,h)}{0}, \underset{(H,\ell)}{0.5}, \underset{(H,h)}{0} \right). \tag{4}$$

These create verification costs of only $(1/6) \cdot 0.5 \cdot c = (1/12)c$.

This improvement is different from the one in example 1(a), as it relies on correlation. The principal benefits from secrecy by exploiting the agent's belief heterogeneity when shifting verification probability from the high to the low signal. This reduces the overall verification probability because the relevant incentive constraint is for type L not to report H, and type L's subjective belief puts more weight on the low signal. Why does this not work in the optimal mechanism? The profitable shift in verification probability when we move from transparency to secrecy requires nonmonotonicity of the total allocation in the signal. However, in any optimal transparent mechanism, the allocation must be pointwise monotone in s for all t. Moving allocation probability toward higher signals for any given type relaxes the upward incentive constraints. As the signal s does not affect the principal's value directly, she is indifferent as to which signals carry allocation probability. This is different in the extension in which s has a direct effect on the principal's value (see app. D).

However, pointwise monotonicity in both type and signal are not enough to obtain transparency. The following example shows that secrecy can also be beneficial when allocation probabilities are interior.

Example 1 continued (c). Consider the total allocation probabilities

$$\boldsymbol{x} + \boldsymbol{z} = \Big(\underset{(L,\ell)}{0.5}, \underset{(L,h)}{0.5}, \underset{(H,\ell)}{1}, \underset{(H,h)}{1}\Big). \tag{5}$$

If the principal has to implement this allocation with a transparent procedure, the cost-minimal verification schedule is

$$\boldsymbol{z} = \Big(\underset{(L,\ell)}{0}, \underset{(L,h)}{0}, \underset{(H,\ell)}{0.5}, \underset{(H,h)}{0.5}\Big).$$

This results in verification costs of $(1/6) \cdot 0.5 \cdot c + (2/6) \cdot 0.5 \cdot c = (1/4)c$. Just as in example 1($b$), the principal can exploit that type L puts more weight on signal ℓ. The cost-minimal verification probabilities under secrecy,

$$\boldsymbol{z} = \Big(\underset{(L,\ell)}{0}, \underset{(L,h)}{0}, \underset{(H,\ell)}{0.75}, \underset{(H,h)}{0}\Big),$$

ensure truthful reporting at lower costs $(1/6) \cdot 0.75 \cdot c = (1/8)c$.

Interior probabilities give the principal more flexibility in shifting verification probability across signals and allow her to benefit from secrecy. In the proof of proposition 1, we show that it is optimal for the principal to use deterministic procedures.

This concludes the intuition for our transparency result under correlation. In general, secrecy enables the principal to exploit the heterogeneous beliefs of agents of different types. However, as the optimal transparent mechanism is monotone in type and signal and nonrandom, there is no room to exploit the channels identified above. For clear illustration, the chosen examples feature fixed allocations, and we present changes in verification only. In the formal proof, instead of decomposing the program, we solve for optimal allocation and verification jointly. This solution strategy is more effective in our setting, as characterizing optimal verification rules for arbitrary allocations is tedious because of the belief heterogeneity.

B. Optimal Transparent Mechanisms

To complete the proof of theorem 2, we need to solve the principal's problem subject to the EPIC constraints.

Proposition 2. The cutoff-with-appeal mechanism presented in theorem 2 maximizes the principal's payoff among all transparently implementable mechanisms.

The formal proof is relegated to appendix A. Solving the problem under EPIC constraints is significantly simpler, as the belief heterogeneity plays no role once the agent knows the signal. In fact, the principal's problem can be solved independently for each signal realization s. Each subproblem corresponding to s is analogous to the degenerate case of a single player in Ben-Porath, Dekel, and Lipman (2014). Hence, proposition 1 creates a link between the cases of correlated and independent information by establishing that the principal voluntarily forgoes the screening potential of heterogeneous beliefs. This is in sharp contrast to settings with money, where the optimal mechanism under independence and the optimal mechanism with the slightest correlation are fundamentally different.

V. Imperfect Verification

Here, we consider imperfect verification. We distinguish two forms of imperfect verification that arise in applications, noisy and coarse verification.

A. Noisy Verification

One cause of imperfect verification in applications is that the technology might occasionally produce incorrect results. For example a low-ability job candidate may sometimes achieve a high test score by luck. As a consequence, verification is *noisy*. This subsection confirms that the key insights from the main part are robust to small noise. In particular, the gain from secrecy is negligible when the noise is small.

Consider a broad class of noisy verification technologies, represented by a family of probabilities $\alpha(\hat{t}|t)$, specifying the probability that true type t can pass as type $\hat{t}$ without being detected in the verification process. We rule out false detections and set $\alpha(t|t) = 1$ for all t. Hence, for $\hat{t} \neq t$, $\alpha(\hat{t}|t) \in [0, 1]$ denotes the verification error, and we denote the largest verification error for a given technology by $\varepsilon = \max_{(\hat{t},t):\hat{t}\neq t}\alpha(\hat{t}|t)$.[19] Given this noisy technology, we present a revelation principle in appendix B

[19] A microfoundation for this reduced-form stochastic verification model, based on pass/fail tests chosen by the principle, can be found in Ball and Kattwinkel (2019). This microfoundation implies additional structure on $\alpha(\hat{t}|t)$.

that allows us to focus on the same class of direct truthful mechanisms defined by $(\boldsymbol{x}, \boldsymbol{z})$.

The "noisy" Bayesian incentive constraints (NBIC) read, for all $t, \hat{t} \in T$,

$$\mathbb{E}_s[x(t,s) + z(t,s)|t] \geq \mathbb{E}_s[x(\hat{t},s) + \alpha(\hat{t}|t)z(\hat{t},s)|t]. \qquad (\mathrm{NBIC}_{t,\hat{t}})$$

When there is noise, with some small probability types can successfully mimic other types, even if the report is verified. Under perfect verification, the optimal mechanism (theorem 2) makes types below the threshold just indifferent between reporting truthfully and lying. This mechanism would not be incentive compatible if verification were noisy: the low types would misreport a type above the threshold, hoping that the noisy verification will not detect this lie. To restore incentive compatibility, type reports below the threshold must be compensated for with a small allocation probability. As a consequence, the optimal transparent mechanism under noisy verification violates the deterministic property discussed in section IV.A.

To illustrate, consider the worst-case technology among all specifications with largest verification error ε. That is, let $\alpha(\hat{t}|t) = \varepsilon$ for all $\hat{t} \neq t$.[20] The optimal transparent mechanism with this uniform-error technology is characterized as follows.

Proposition 3. Suppose that the verification technology has uniform error probability $\varepsilon \in (0,1)$. Then, the optimal transparent mechanism is as follows: there are cutoffs $\underline{s}^\varepsilon < \overline{s}^\varepsilon$ and an appeal threshold $\overline{t}^\varepsilon$ such that

(i) if $s \geq \overline{s}^\varepsilon$, then $x(t,s) = 1$ for all t and $z(t,s) = 0$ for all t;

(ii) if $s \in [\underline{s}^\varepsilon, \overline{s}^\varepsilon)$,

then $x(t,s) = \begin{cases} 0 & \text{for } t \geq \overline{t}^\varepsilon \\ \varepsilon & \text{for } t < \overline{t}^\varepsilon \end{cases}$ and $z(t,s) = \begin{cases} 1 & \text{for } t \geq \overline{t}^\varepsilon, \\ 0 & \text{for } t < \overline{t}^\varepsilon; \end{cases}$

(iii) if $s < \underline{s}^\varepsilon$, then $x(t,s) = 0$ for all t and $z(t,s) = 0$ for all t.

In all three cases,

$$\overline{t}^\varepsilon = \min\left\{t \middle| v(t) - \frac{c}{1-\varepsilon} > 0\right\},$$

$$\overline{s}^\varepsilon = \min\left\{s \middle| \mathbb{E}_t[v(t)|s] - \mathbb{E}_t\left[\left(v(t) - \frac{c}{1-\varepsilon}\right)^+ \middle| s\right] > 0\right\}, \text{ and}$$

$$\underline{s}^\varepsilon = \min\left\{s \middle| \mathbb{E}_t[v(t)|s] + \frac{1-\varepsilon}{\varepsilon}\mathbb{E}_t\left[\left(v(t) - \frac{c}{1-\varepsilon}\right)^+ \middle| s\right] > 0\right\}.$$

[20] This is the worst case for the principal: consider any mechanism that is incentive compatible when $\alpha(\hat{t}|t) = \varepsilon$ for all $\hat{t} \neq t$. This mechanism is also incentive compatible under any verification technology for which the largest verification error is ε.

Figure 2 depicts a mechanism defined in the result above. To see the differences with the optimal mechanism under perfect verification (theorem 2), consider a signal between the two cutoffs, $s \in (\underline{s}_\varepsilon, \bar{s}_\varepsilon)$: in this case, reports above the threshold $\bar{t}^\varepsilon$ are verified and get the good. To make sure that types below the threshold do not misreport and hope to get the good with probability ε after (failed) verification, the mechanism gives allocation probability ε for type reports below the threshold $\bar{t}^\varepsilon$. Given the s, the principal's expected payoff from this additional allocation probability is $\varepsilon\mathbb{E}[v(t)\mathbf{1}_{\{t<\bar{t}^\varepsilon\}}|s]$, which is negative for $s < \bar{s}^\varepsilon$. Thus, the noise causes an additional, indirect cost of verification. Therefore, the appeal threshold $\bar{t}^\varepsilon$ must be higher than with perfect verification. The above result shows that the threshold with error ε is equal to the optimal threshold under perfect verification if the cost c were increased to $c/(1-\varepsilon)$. The same is true for the upper signal cutoff $\bar{s}^\varepsilon$. For low signals s, the expected costs of verification may now outweigh the expected benefit from allocating to types $t \geq \bar{t}^\varepsilon$.

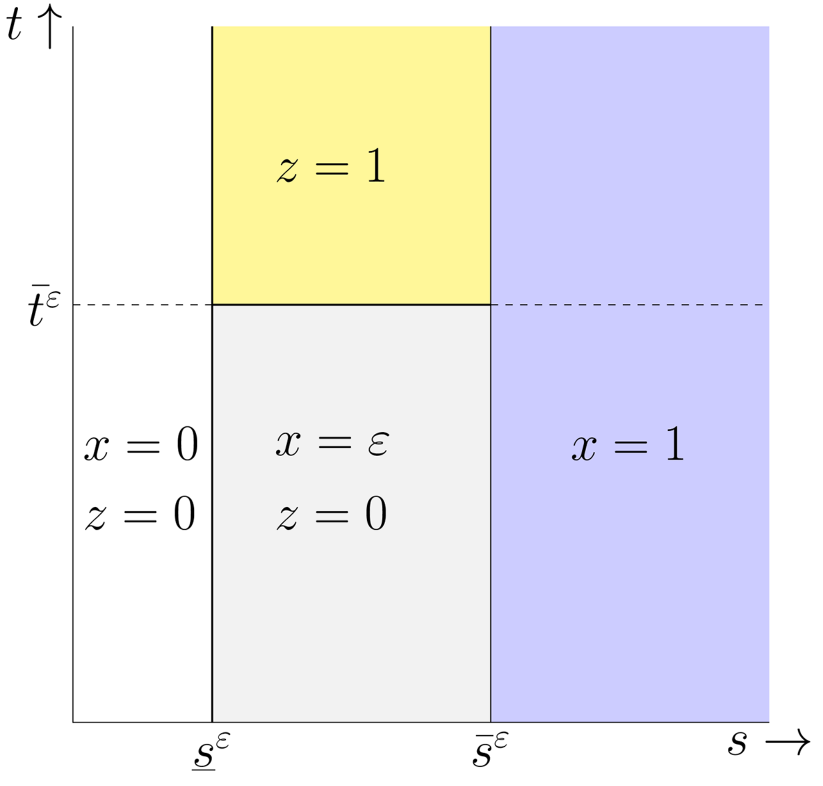

FIG. 2.—Noisy verification.

This leads to another difference with respect to the optimal mechanism under perfect verification: there is an additional cutoff $\underline{s}^{\varepsilon}$ such that for signals $s < \underline{s}^{\varepsilon}$, the mechanism does not allocate or verify after any report. Note that the lower cutoff $\underline{s}^{\varepsilon}$ is increasing in the error probability ε; in particular, if ε is close enough to 0, then $\underline{s}^{\varepsilon} = \min S$ and case (iii) in proposition 3 is never met.

The optimal mechanism above, together with the illustrations in section IV.A, already indicates that the gain from secrecy under noisy verification is not zero. For $\varepsilon \in (0, 1)$, the optimal transparent mechanism is not deterministic, violating one of the three properties that rule out gains from secrecy. On the basis of example 1(*c*), it is easy to construct an improvement for the principal that exploits secrecy by shifting the allocation probabilities ε across signals to exploit the agent's belief heterogeneity.

However, this modification and the resulting gain from secrecy are only of the order of magnitude of the noise in verification. In particular, our key results are robust to small noise, in that the gain from secrecy vanishes linearly in the verification error ε.

Proposition 4. Let ε be the largest verification error. The gain from secrecy is bounded by

$$\varepsilon \cdot \mathbb{E}_t[-v(t)\mathbf{1}_{\{t<\bar{t}\}}\mathbf{1}_{\{s<\bar{s}\}}], \tag{6}$$

where $\bar{s}$ and $\bar{t}$ are defined in theorem 2.

Transparent mechanisms have many advantages (as discussed in sec. III). This proposition guarantees that the loss for the principal from using a transparent procedure is small when the verification error is small. Proposition 4 is thus a generalization of our transparency result, proposition 2.[21]

The gains from secrecy under noisy verification arise from exploiting the agent's belief heterogeneity to shift the (interior) compensation probabilities ε across signals. The gain from these shifts can therefore be bounded by equation (6), which measures exactly the indirect cost the principal incurs from these compensation probabilities.

The proof of proposition 4 in appendix A establishes the following, tighter, bound:

$$\varepsilon \cdot \sum_{s<\bar{s}} \mathbb{P}[s] \min\left\{\frac{1}{\varepsilon}\mathbb{E}_t[(v(t) - c)^+|s]; \mathbb{E}_t[-v(t)\mathbf{1}_{\{t<\bar{t}\}}|s]\right\}.$$

The minimum is equal to its first term when the indirect cost of allocating with probability ε to types below the threshold exceeds the benefit from verifying and allocating above the threshold (case (iii) in proposition 3).

[21] Note that the definition of $\bar{s}$ implies that the bound (6) is positive.

The uniform error is a worst-case scenario. The optimal transparent mechanism in proposition 3 is also incentive compatible in any other scenario with maximal verification error ε. That is why the bound in proposition 4 holds across all specifications of noisy verification with maximal error probability ε.

B. *Coarse Verification*

Another form of imperfection arises in applications when verification is implemented through the presentation of hard evidence: depending on his type, the agent has different pieces of evidence at his disposal, which he can present to the principal. The perfect verification from the main part would correspond to an evidence structure in which for every type there is some evidence that only this type can present. Oftentimes, however, some different types have access to the same pieces of evidence. For example, the possible ability levels of a job candidate are usually more complex than the available degrees and certificates that can be demanded to verify his claims. As a consequence, verification in these cases is *coarse*, such that some groups of types remain indistinguishable in the verification process.

We incorporate this form of imperfection, assuming that some neighboring types are indistinguishable. Formally, consider a partition $\mathcal{T}$ of the type space T into discrete intervals τ containing successive types who hold the same pieces of evidence, that is, $\mathcal{T} = \{\tau_1, \dots, \tau_n\}$, where for any $t_i < t_j < t_k \in T$, if $t_i, t_k \in \tau$, then also $t_j \in \tau$. Denote by $\tau(t)$ the interval in the partition that contains type t. If the principal verifies and the agent's type is t, the principal learns $\tau(t)$.[22] The principal's problem has the same objective as the original (LP). The "coarse Bayesian" IC constraints (CBIC) are as follows: for all $t, \hat{t}$ with $\hat{t} \notin \tau(t)$, ($\text{BIC}_{t,\hat{t}}$), and, additionally, for all $t, \hat{t}$ with $\hat{t} \in \tau(t)$:

$$\mathbb{E}_s[x(t,s) + z(t,s)|t] \geq \mathbb{E}_s[x(\hat{t},s) + z(\hat{t},s)|t]. \qquad (\text{CBIC}_{t,\hat{t}})$$

If type t reports a type $\hat{t} \in \tau(t)$ with the same evidence, t receives the good with the same probabilities as $\hat{t}$, independent of whether report $\hat{t}$ is verified or not. Clearly, constraints ($\text{CBIC}_{t,\hat{t}}$) are more demanding than constraints ($\text{BIC}_{t,\hat{t}}$). Nevertheless, the optimal mechanism can still be implemented transparently.

[22] We could modify this setting and assume that each type, in addition, has all the evidence pieces of any lower type. Formally, $\tau'(t) = \cup_{\tilde{t} \leq t} \tau(\tilde{t})$. This formulation captures the idea that a more capable job candidate can produce all certificates a less capable candidate can produce. Our results remain unchanged under this specification.

See Bull and Watson (2004) and Green and Laffont (1986) for general treatments of verifiable evidence where access differs across types. Our specification satisfies the normality (or nested-range) condition, which these papers show to be sufficient to focus on truthful direct mechanisms.

PROPOSITION 5. The principal's problem under coarse verification is solved by the cutoff mechanism with signal cutoffs s_c^*, $\bar{s}_c$ and type threshold $\bar{t}$ such that

(i) if $s \geq \bar{s}_c$, then $x(t, s) = 1$ for all t and $z(t, s) = 0$ for all t;
(ii) if $s_c^* \leq s < \bar{s}_c$, then $x(t, s) = 0$ for all t and

$$z(t, s) = \begin{cases} 1 & \text{if } t \geq \min(\tau(\bar{t})), \\ 0 & \text{otherwise;} \end{cases}$$

(iii) if $s < \min\{s_c^*, \bar{s}_c\}$, then $x(t, s) = 0$ for all t and

$$z(t, s) = \begin{cases} 1 & \text{if } t > \max(\tau(\bar{t})), \\ 0 & \text{otherwise,} \end{cases}$$

where $\bar{t} = \min\{t | v(t) - c > 0\}$, as in theorem 2, and $\bar{s}_c$ and s_c^* are defined as follows:

$$\bar{s}_c = \min\left\{s | \mathbb{E}_t[v(t)|s] > \mathbb{E}_t[(v(t) - c)\mathbf{1}_{\{t > \max(\tau(\bar{t}))\}}|s] + (\mathbb{E}_t[(v(t) - c)\mathbf{1}_{\{t \in \tau(\bar{t})\}}|s])^+\right\}, \text{ and}$$

$$s_c^* = \min\left\{s | \mathbb{E}_t[(v(t) - c)\mathbf{1}_{\{t \in \tau(\bar{t})\}}|s] > 0\right\}.$$

This mechanism is transparent.

Figure 3 sketches an optimal mechanism under coarse verification for the case $\bar{s}_c > s_c^*$. If the signal realization is above the cutoff $\bar{s}_c$, the principal allocates the good to the agent irrespective of his reported type and without verification ($x = 1$). If the signal is below the cutoff $\bar{s}_c$, the agent can receive the good only after his type is verified ($z = 1$) to be high enough. For $s \in [s_c^*, \bar{s}_c)$, all types in $\tau(\bar{t})$ and in all intervals above receive the good after verification. For $s < s_c^*$, only types in the intervals strictly above $\tau(\bar{t})$ receive the good after verification.

Note that in the optimal mechanism under perfect verification (theorem 2), the verification threshold is independent of the signal; types above $\bar{t}$ are verified, as their value $v(t)$ is above c. In contrast, coarse verification reveals only the agent's partition interval. Within the interval $\tau(\bar{t})$, $v(t)$ may be below or above c. The information about t contained in the signal affects the conditional expectation of $v(t)$. Hence, types in $\tau(\bar{t})$ are verified only if the expected value of allocating to all types within $\tau(\bar{t})$ (conditional on s) is worth paying the verification cost c. As a consequence, the mechanism has an additional signal cutoff s_c^*.[23]

Proposition 5 confirms that this mechanism is transparent. One can easily verify that the three properties introduced above—monotonicity

[23] Note that when $\bar{t} = \min(\tau(\bar{t}))$, we have $s_c^* = \min(S)$, and the optimal mechanism characterized in proposition 5 is the same as the optimal mechanism under perfect verification. Intuitively, whenever the (coarse) verification technology allows the principal to distinguish whether the type is below the optimal appeal threshold $\bar{t}$ or not, the principal can implement the optimal mechanism from theorem 2. Whenever feasible, this must be optimal in the more restrictive principal's problem under coarse verification.

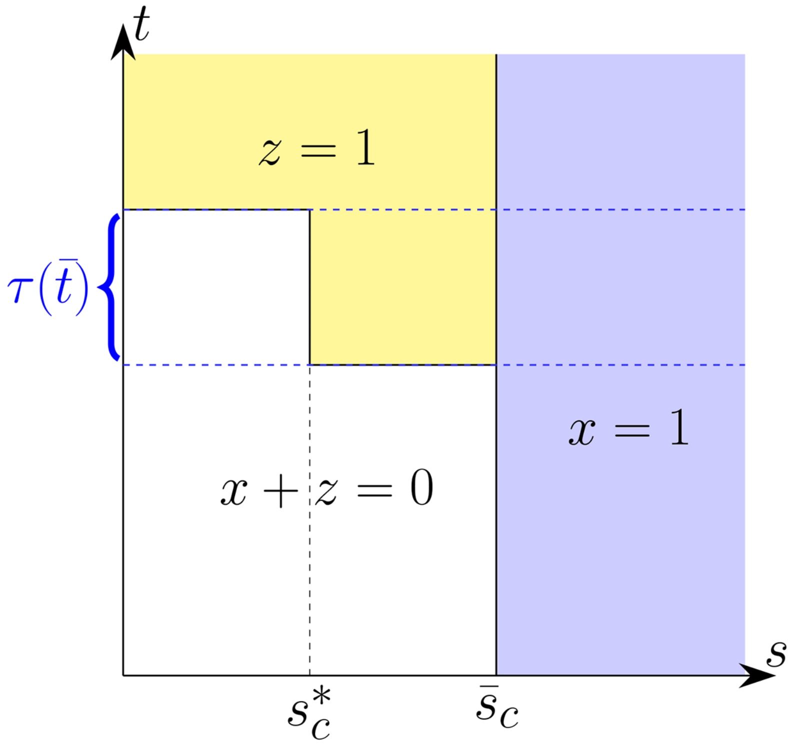

FIG. 3.—Coarse verification.

in type, monotonicity in signal, and nonrandomness—are satisfied. There is no loss from transparency.

VI. Literature

A. *Mechanism Design with Correlation*

The role of correlated information in mechanism design with monetary transfers was first studied in Myerson (1981) and Crémer and McLean (1988).[24] In these papers, the principal exploits correlation to screen the agent's type through his heterogeneous beliefs. With unbounded monetary transfers, this allows the principal to elicit the agent's private information without paying him any information rent. On rent-free elicitation of all private information, Neeman (2004, 56) notes, "This implication has made several economists uncomfortable," as it implies that the agent's

[24] See also Riordan and Sappington (1988), Johnson, Pratt, and Zeckhauser (1990), and McAfee and Reny (1992).

private information is irrelevant. We present a setting in which rent-free information elicitation is not possible despite the "beliefs-determine-preferences" (Neeman 2004) property being fulfilled. More than that, in our optimal mechanism, the principal abstains from exploiting the belief heterogeneity altogether by making her signal public. Demougin and Garvie (1991) show that the qualitative insights from Crémer and McLean (1988) also apply in the case of bounded transfers or limited liability. Therefore, different from our setting, the principal gains from keeping her signal secret.[25]

Milgrom and Weber (1982) and Eső and Szentes (2007) also study different informational regimes in mechanism design and find that information revelation can benefit the principal. However, the question about transparency we seek to address is distinct from the comparisons in these papers. We ask whether the principal profits from concealing her available private information from the agent. They ask whether the principal profits from creating additional information that is observed by the agents and can otherwise not be used by the principal.[26] In these settings, raising more information entails a trade-off between increasing efficiency and increasing information rents. In our setting, the principal can always observe and condition the mechanism on her signal so a more informative signal is always beneficial for her, and it is immediate that she weakly prefers concealing her information. Our main result shows that she does not strictly prefer concealing. Our perspective is parallel to Jehiel's (2015, 736) question: "When is it best for the Principal to commit to not disclosing all that he/she knows in moral hazard interactions?" In contrast, we analyze an adverse-selection problem and find that secrecy is never strictly beneficial.

B. *Costly State Verification*

Costly state verification, the possibility for the mechanism designer to learn the agent's private information at a cost, was first introduced by Townsend (1979) to a principal-agent model of optimal debt contracts.[27] More recently, the verification technology proposed by Townsend (1979) has been analyzed in models without transfers by Ben-Porath, Dekel, and

[25] In the absence of monetary transfers, Bhargava, Majumdar, and Sen (2015) show that positively correlated beliefs among voters eliminate the impossibility of nondictatorial voting rules established by Gibbard (1973) and Satterthwaite (1975).

[26] In Eső and Szentes (2007), the additional information is privately observed by the agents. In Milgrom and Weber (1982), the auction mechanism is exogenously fixed; in sec. 7 of their paper, the authors allow the principal to set a reserve price contingent on the additional information, but only when the information is also revealed to the agents. McAfee and Reny (1992) show that if the principal could design any mechanism, not revealing the signal to the agents would allow her to extract all surplus, outperforming the auctions discussed in Milgrom and Weber (1982).

[27] See also Gale and Hellwig (1985) and Mookherjee and Png (1989).

Lipman (2014), Beshkar and Bond (2017), Erlanson and Kleiner (2020), Halac and Yared (2020), and Li (2020).[28] Our model is most closely related to that of Ben-Porath, Dekel, and Lipman (2014), who study an allocation problem with multiple agents. Our main departure from this literature lies in the correlated information. Just as with monetary transfers, the belief heterogeneity induced by correlation impedes the standard approach to characterize mechanisms through interim expectations. New techniques are required. Perhaps surprisingly—and unlike with monetary transfers—optimal mechanisms share the qualitative features discovered in the above papers where correlated information is absent: a simple cutoff structure is optimal, and mechanisms are EPIC; that is, the agent would also report truthfully if he were informed about the signal before his report.[29]

The equivalence between BIC and EPIC mechanisms in verification settings under independent information holds more generally rather than only for optimal mechanisms (see Erlanson and Kleiner 2020).[30] Our paper illustrates the importance of the independence assumption for this question. Correlation creates an additional potential benefit of secrecy in BIC mechanisms. We establish conditions that determine whether this benefit for the designer exists. We discuss the implications of our results for court procedures. Silva (2019) and Siegel and Strulovici (2020) also consider the role of verification in this context.

As the principal has private information, our model is also related to the informed-principal problem; see Myerson (1983) and Maskin and Tirole (1990). A priori, the assumption that the principal's signal is contractible sets us apart from this strand of literature, as it allows the principal to commit to a mechanism before learning the signal. Thus, the proposed mechanism does not convey any information to the agent. However, since the optimal mechanism we derive is EPIC, it also constitutes a solution to the informed-principal game in which the principal proposes a mechanism only after observing her private signal: this game has a separating equilibrium in which the principal proposes one mechanism for each signal realization and, thus, the agent perfectly learns the signal from the

[28] Mechanism design with alternative verification technologies or evidence is studied in Green and Laffont (1986), Bull and Watson (2004), Glazer and Rubinstein (2004), Deneckere and Severinov (2008), Kartik and Tercieux (2012), Hart, Kremer, and Perry (2017), Mylovanov and Zapechelnyuk (2017), Ball and Kattwinkel (2019), Ben-Porath, Dekel, and Lipman (2019), Epitropou and Vohra (2019), and Koessler and Perez-Richet (2019), as well as Rappoport (2020). More recently, Silva (2021) studies a model with an analogous informational setup where the principal's verification is costless but yields imperfect results. This relates to our extension in sec. V.A.

[29] Allocation mechanisms without money or verification are the focus of Börgers and Postl (2009), Goldlücke and Tröger (2018), and Ortoleva, Safonov, and Yariv (2021) under independence and of Kattwinkel (2020) under correlation.

[30] This relates to Manelli and Vincent (2010) and Gershkov et al. (2013), who show equivalence between BIC and dominant-strategy IC mechanisms in settings with monetary transfers and independent information.

proposal.[31] With monetary transfers, Cella (2008) and Severinov (2008) show that correlated information allows for an efficient solution to the informed-principal problem. Skreta (2011) analyzes when an informed principal can generate exploitable correlation between the agents by partial information disclosure. In our setting, without monetary transfers, the informed principal cannot exploit the secrecy of her signal. In Ottaviani and Prat (2001), an informed monopolist is assumed to propose a menu before eliciting the buyer's type. In this restricted informed-principal problem, the monopolist benefits from revealing her private information to the buyer. They note that under the classical informed-principal approach, their monopolist would pull a Crémer-McLean and extract all the surplus by keeping her information secret.

VII. Conclusion

This paper studies the role of correlated information in a mechanism-design model in which the principal may use costly verification instead of monetary transfers to incentivize the revelation of private information. We show that a transparent mechanism is optimal. It is without loss for the principal to make her information public before contracting with the agent. Our result gives a rationale for the use of transparent procedures in a variety of applications from hiring to procedural law. This is in contrast with results on correlation in mechanism-design problems with money.

We confirm that our results are robust to small imperfections in the verification technology.

In the online appendix, we present two extensions of our model. First, we consider the case in which the agent also bears some cost of being verified. Again, we show that our findings are robust to small agent costs. Second, we allow for the principal's signal to have a direct effect on her payoffs. When this effect is positive, our result carries over: transparency comes without loss for the principal. When, in contrast, the direct effect is negative, the principal benefits from secrecy. This reveals that correlated information poses a limitation to the general equivalence between BIC and EPIC mechanisms.

Appendix A

Omitted Proofs

A1. Proof of Theorem 1

The proof of theorem 1 consists of a revelation principal (lemma 1) and optimality arguments (lemma 2) to further reduce the class of mechanisms.

[31] In this separating equilibrium, the mechanisms proposed by the principal are indexed by the signal realizations, although different realizations may lead to the same mechanism. Alternatively, there is an equilibrium in which the principal proposes only two mechanisms, for signals below and above the cutoff.

The principal faces the fundamentals (X, T, S, f, u), where $X = \Delta(\{0, 1\})$ denotes the feasible outcomes (allocation probabilities) and T and S are type and signal spaces with joint distribution f. The marginals of f are denoted by $f_S \in \Delta(S)$ and $f_T \in \Delta(T)$. Finally, u is the agent's utility function. Given these fundamentals, the principal designs an extensive game form $G = \langle H, \prec, A, \mathcal{A}, P, \delta_C, \mathcal{I}, g \rangle$. We borrow the formal definition from Akbarpour and Li (2020) and extend it to allow for verification stages and chance moves as part of the game. The player function P, with $P(h) \in \{1, V, C\}$, states for every nonterminal history h whether the agent (player 1) is called to play, verification happens (player V), or chance (player C) is called to play. If $P(h) = 1$, the agent chooses an action from $A(I)$, where information set I is the element of the partition $\mathcal{I}$ of $\{h' : P(h') = 1\}$, which contains h. If $P(h) = V$, verification occurs. This means that the game's structure from this history onward may depend on the real agent type t.[32] If $P(h) = C$, the successor of history h, that is, $h' \succcurlyeq h$, is chosen according to the predetermined probability measure $\delta_C(h)$. Since the principal has commitment power when designing the game, it is without loss to let the game start with the draw of the signal. That is, $P(h_\varnothing) = C$ and $\delta_C(h_\varnothing) = f_S$.

Lemma 1 (revelation principle). For any pair (G, σ) of extensive-form game G and equilibrium strategy σ, there is an incentive-compatible and outcome-equivalent direct mechanism of the following form: there are functions $e : T \times S \to [0, 1]$ and $a : T \times S \times T \cup \{\varnothing\} \to [0, 1]$.

1. The agent reports type $\hat{t} \in T$.
2. Given report $\hat{t}$ and signal s, verification occurs with probability $e(\hat{t}, s)$ and reveals the agent's real type, t.
3. Allocation probability is $a(\hat{t}, s, t)$ if verification took place and $a(\hat{t}, s, \varnothing)$ if no verification took place.

Proof. Consider game G together with a collection of strategies $\sigma = (\sigma_t)_{t \in T}$, where for each t, strategy $\sigma_t : \mathcal{I} \to A$, such that $\sigma_t(I) \in A(I)$ for all $I \in \mathcal{I}$ and σ_t is a best response.[33] The following proof reduces the pair (G, σ) to a direct mechanism of the above class such that (i) each type prefers to report truthfully and (ii) given truthful reporting, the verification probabilities and expected allocation outcomes resulting from the direct mechanism are equivalent to the ones in the original game.

After report $\hat{t}$ and signal realization s, the direct mechanism specifies verification probability $e(\hat{t}, s)$. To determine $e(\hat{t}, s)$, the principal simulates the play of game G from signal realization s onward, assuming that the agent plays strategy

[32] This verification technology assumes that verification results in exact knowledge of the agent's type. Alternatively, we may define verification as a technology that requires a type report from the agent and reveals whether this report and the true type coincide. All results in the paper are unaffected by this choice.

[33] The restriction to pure strategies is without loss for our purposes. Any pair of single-player game and equilibrium strategy in mixed strategies can be equivalently designed as a game of pure strategies where the agent's mixed action is replaced by a chance move that simulates the mixing.

$\sigma_{\hat{t}}$, and chooses $e(\hat{t}, s)$ equal to the probability that the simulated equilibrium path entails at least one verification stage.[34]

Subsequently, the allocation is chosen depending on whether the result of the verification is "no verification" ($\varnothing$), "truthful report" ($\hat{t}$), or "misreport" ($t \notin \{\varnothing, \hat{t}\}$). In the first case, $a(\hat{t}, s, \varnothing)$ is chosen equal to the expected allocation probability that results from the simulated play, conditional on the event that no verification stage was reached. In the second case, $a(\hat{t}, s, \hat{t})$ is determined by simulating play conditional on verification with strategy $\sigma_{\hat{t}}$ and choosing the continuation game after any verification stage that corresponds to real type $\hat{t}$. Finally, in the case of a detected misreport, $a(\hat{t}, s, t)$ is determined as follows. The principal simulates play, conditional on verification, with strategy $\sigma_{\hat{t}}$ until the first verification stage. After this stage, play of the continuation game corresponding to t is simulated with strategy σ_t. Strategy σ_t may rule out reaching the above path. However, as any strategy specifies action choices for all information sets of the game, probability $a(\hat{t}, s, t)$ is well defined.

This construction (i) ensures truthful reporting: when t is reported truthfully in the direct mechanism, the agent's expected payoff is equal to the one resulting from strategy σ_t in the original game. Furthermore, the outcome resulting from a misreport $\hat{t} \neq t$ in the direct mechanism was available to type t in the original game by playing strategy $\sigma_{\hat{t}}$ until the first verification (if any) and playing according to σ_t afterward. Furthermore, it follows from the construction that (ii) the verification and allocation probabilities, and therefore all expected payoffs, in the original game equilibrium and the direct mechanism coincide. QED

Corollary 1 (information design). The extensive game form allows fully flexible information design by designing chance moves that depend on signal realization s and the design of information sets for the agent. By the revelation principle, any benefit from information design can be achieved within the reduced class of direct mechanisms.

To exemplify this, recall that the first stage of the game was chosen to be the signal realization. The flexibility to design the rest of the game G with arbitrary chance moves and information sets for the agent allows the principal to freely design any communication about the signal. For example, assuming that the agent is called to play after the signal draw (history h_s), specifying $I(h_s) = h_s$ for all s represents the situation where the agent learns the signal realization directly, whereas $I(h_s) = \{h_s\}_{s \in S}$ gives no information to the agent. More complex, the principal may design different chance moves $\delta_C(h_s)$ for each realization s and collect the following histories into information sets to achieve any Bayes-plausible interim belief distribution.

The next observation uses the principal's objective to further restrict the set of direct mechanisms, ruling out some strictly suboptimal mechanisms.

Lemma 2 (maximal punishment and minimal verification). It is without loss of optimality for the principal to restrict the class of direct mechanisms to fulfill the following properties:

1. maximal punishment: if $t \notin \{\varnothing, \hat{t}\}$, then $a(\hat{t}, s, t) = 0$; and
2. minimal verification: if $e(t, s) > 0$, then $a(t, s, t) = 1$.

[34] Reaching multiple verification stages will not be part of a principal-optimal game, but for the revelation principle is formulated for general games.

Proof. For the first property, note that detecting a misreport $t \notin \{\varnothing, \hat{t}\}$ is an off-path event in the direct mechanism, so that modifying the allocation probability has no effect on the principal's payoff. For the agent, setting the allocation after a misreport as low as possible decreases incentives to misreport while maintaining the payoff of any truth-telling type unchanged.

Second, suppose that $e(t, s) > 0$ and $a(t, s, t) < 1$. One could now lower the probability of verification $e(t, s) \downarrow$ while increasing the probability of allocation after confirming the report as true $a(s, t, t) \uparrow$ such that $e(t, s)a(t, s, 1)$ remains constant.

Lowering the verification probability would increase the incentives to misreport and the overall allocation probability after report t and signal s only if there was allocation with positive probability conditional on no verification, that is, $a(t, s, \varnothing) > 0$. However, in this case, this allocation could be lowered $a(t, s, \varnothing) \downarrow$ such that $(1 - e(t, s))a(t, s, \varnothing)$ remains constant. Then, the incentives to misreport and the overall allocation probability would remain constant. As these procedures would save verification costs while keeping all unconditional allocation probabilities constant, the fact that an optimal mechanism features nonmaximal reward can be ruled out. QED

To conclude the proof of theorem 1, define by $z(t, s) = e(t, s)$ the joint probability of verification and allocation and by $x(t, s) = (1 - e(t, s))a(t, s, \varnothing)$ the joint probability of no verification and allocation. The set of mechanisms described by

$$\{(x(t, s), z(t, s))_{t \in T, s \in S} | \ \forall\ t \in T \ \ \forall\ s \in S : 0 \ \leq \ x(t, s) + z(t, s) \ \leq \ 1\}$$

is equivalent to all minimal verification and maximal punishment direct mechanisms.[35] QED

A2. *Proof of Proposition 1*

We present the following relaxation of the problem and show that it is solved by a transparent mechanism that is also feasible in the original (LP). Define the set of profitable types as those t with a positive allocation value,

$$T^{+} \equiv \{t \in T | v(t) > 0\},$$

and the unprofitable types accordingly as $T^{-} \equiv T \backslash T^{+}$. Both sets are nonempty by the assumption that v crosses 0. Otherwise, the optimal mechanism is trivial. The relaxed problem includes only those incentive constraints that prevent types in T^{-} from misreporting types in T^{+}. Hence, it reads as follows:

[35] The inverse mapping is given by

$$(e(t, s), a(t, s, \varnothing)) = (z(t, s), x(t, s)/(1 - z(t, s))).$$

Note that the value of $a(t, s, \varnothing)$ does not play any role in the mechanism if $e(t, s) = z(t, s) = 1$ and can therefore be chosen arbitrarily.

$$\max_{(\boldsymbol{x},\boldsymbol{z})\geq 0} \sum_{t\in T}\sum_{s\in S} f(t,s)[x(t,s)v(t) + z(t,s)(v(t)-c)],$$

$$\text{subject to } \forall\ (t,\hat{t}) \in T^- \times T^+ : (\text{BIC}_{t,\hat{t}}) \quad \text{and} \qquad \text{(LP.r)}$$

$$\forall\ (t,s) \in T \times S:\ x(t,s) + z(t,s) \leq 1.$$

In the remainder of the proof, we derive feasible changes to a solution to the relaxed problem that do not lower the principal's value and that finally lead to a transparent cutoff mechanism. We make repeated use of the following notation: we denote changes in the allocation probability by $dx(t,s)$, so that the new probability after the change is given by $x(t,s) + dx(t,s)$, where $dx(t,s)$ may be positive or negative, and analogously for $dz(t,s)$. Further, $d(\text{BIC}_{t,\hat{t}})$ denotes the change in surplus utility that type t receives from reporting the truth rather than misreporting $\hat{t}$, which is induced by a change of the above form. Recall that the constraint $(\text{BIC}_{t,\hat{t}})$ reads as $\Sigma_s f(t,s)[x(t,s) + z(t,s) - x(\hat{t},s)] \geq 0$, so that $d(\text{BIC}_{t,\hat{t}})$ denotes the change to the left-hand side of this inequality.

The value for the principal is given by

$$V = \sum_{t\in T}\sum_{s\in S} f(t,s)[x(t,s)v(t) + z(t,s)(v(t)-c)],$$

and dV will denote the induced change to this value.

A2.1. Step 1

The optimal mechanism in the relaxed problem features $\forall\ t \in T^-\ \forall\ s \in S: z(t,s) = 0$.

Suppose that $z(t,s) > 0$ for some type $t \in T^-$. Shifting probability mass from $z(t,s)$ to $x(t,s)$ such that the overall allocation probability stays constant,

$$0 < dx(t,s) = -dz(t,s),$$

saves the principal verification costs and does not distort the incentives, as type t's incentive to misreport remains the same and all incentive constraints to misreport a type $t \in T^-$ are ignored in the relaxed problem.

A2.2. Step 2

There is an optimal mechanism in the relaxed problem featuring a cutoff form for $\boldsymbol{x}(\hat{t},\cdot)$:

$$\forall\ \hat{t} \in T^+\ \exists\ \tilde{s}(\hat{t}) \in S: x(\hat{t},s) \begin{cases} = 0 & \text{if } s < \tilde{s}(\hat{t}), \\ \in [0,1) & \text{if } s = \tilde{s}(\hat{t}), \\ = 1 & \text{if } s > \tilde{s}(\hat{t}). \end{cases}$$

Take a feasible IC mechanism of the relaxed problem featuring that for some $\hat{t} \in T^+$, $\exists\ s < s' \in S$ such that $x(\hat{t},s) > 0$ and $x(\hat{t},s') < 1$.

Modify the mechanism only at two points, shifting allocation probability mass from $x(\hat{t},s)$ to $x(\hat{t},s')$; that is, $dx(\hat{t},s) < 0$ and $dx(\hat{t},s') > 0$. Choose these shifts in a proportion such that for the highest unprofitable type, $\tilde{t} \equiv \max T^-$, the incentive to misreport $\hat{t}$ remains unchanged:

$$0 \overset{!}{=} d(\mathrm{BIC}_{\tilde{t},\hat{t}}) = -f(\tilde{t}, s)dx(\hat{t}, s) - f(\tilde{t}, s')dx(\hat{t}, s')$$
$$= 0 \Leftrightarrow dx(\hat{t}, s) = -\frac{f(\tilde{t}, s')}{f(\tilde{t}, s)}dx(\hat{t}, s').$$

For all types $t \in T^-$, we have $t \leq \tilde{t}$, and, therefore,

$$d(\mathrm{BIC}_{t,\hat{t}}) = -f(t, s)dx(\hat{t}, s) - f(t, s')dx(\hat{t}, s')$$
$$= f(t, s)\left[\frac{f(\tilde{t}, s')}{f(\tilde{t}, s)} - \frac{f(t, s')}{f(t, s)}\right]dx(\hat{t}, s') \geq 0$$

by the MLRP. The principal's value changes in the following way:

$$dV = f(\hat{t}, s)dx(\hat{t}, s)v(\hat{t}) + f(\hat{t}, s')dx(\hat{t}, s')v(\hat{t})$$
$$= f(\hat{t}, s)\left[-\frac{f(\tilde{t}, s')}{f(\tilde{t}, s)}dx(\hat{t}, s')\right]v(\hat{t}) + f(\hat{t}, s')dx(\hat{t}, s')v(\hat{t})$$
$$= f(\hat{t}, s)\left[\frac{f(\hat{t}, s')}{f(\hat{t}, s)} - \frac{f(\tilde{t}, s')}{f(\tilde{t}, s)}\right]dx(\hat{t}, s')v(\hat{t}) \geq 0,$$

since $dx(\hat{t}, s') > 0$ and $\hat{t} \in T^+$, which implies both $v(\hat{t}) \geq 0$ and $\hat{t} > \tilde{t}$. The proposed shift is clearly feasible if in the original mechanism, $x(\hat{t}, s') + z(\hat{t}, s') < 1$. In the case that $x(\hat{t}, s') + z(\hat{t}, s') = 1$, it can still be implemented by shifting in addition mass from $z(\hat{t}, s')$ to $z(\hat{t}, s)$ to ensure that $x(\hat{t}, s') + z(\hat{t}, s')$ and $x(\hat{t}, s) + z(\hat{t}, s)$ remain constant:

$$dx(\hat{t}, s') + dz(\hat{t}, s') = 0 \quad \text{and} \quad dx(\hat{t}, s) + dz(\hat{t}, s) = 0.$$

This implies that $dz(\hat{t}, s') < 0$ and $dz(\hat{t}, s) > 0$. This is feasible, as $x(\hat{t}, s') < 1$ and $x(\hat{t}, s') + z(\hat{t}, s') = 1$ imply that $z(\hat{t}, s') > 0$. As $x(\hat{t}, s) > 0$, we must further have $z(\hat{t}, s) < 1$ by feasibility. To maintain the total allocation probabilities constant, the above changes in x are compensated for by the following changes in z:

$$dz(\hat{t}, s) = \frac{f(\tilde{t}, s')}{f(\tilde{t}, s)}(-dz(\hat{t}, s')).$$

The incentives for any lower type to misreport his type as $\hat{t}$ are weakened in the same way as above because $z(\hat{t}, s)$ and $z(\hat{t}, s')$ do not play a role in the constraints that prevent misreport $\hat{t}$.

Finally, the principal's value now changes by

$$dV = f(\hat{t}, s)[dx(\hat{t}, s)v(\hat{t}) + dz(\hat{t}, s)(v(\hat{t}) - c)] + f(\hat{t}, s')$$
$$[dx(\hat{t}, s')v(\hat{t}) + dz(\hat{t}, s')(v(\hat{t}) - c)]$$
$$= -c[f(\hat{t}, s)dz(\hat{t}, s) + f(\hat{t}, s')dz(\hat{t}, s')]$$
$$= -cf(\hat{t}, s)\left[\frac{f(\tilde{t}, s')}{f(\tilde{t}, s)} - \frac{f(\hat{t}, s')}{f(\hat{t}, s)}\right](-dz(\hat{t}, s')) \geq 0,$$

as, by the MLRP, the term in square brackets in the last line is negative and, by assumption, $-\ dz(\hat{t}, s') \geq 0$.

A2.3. Step 3

There is an optimal mechanism in the relaxed problem featuring $\boldsymbol{x}(\hat{t},\cdot) = \boldsymbol{x}(\hat{\hat{t}},\cdot)$ for all $\hat{t}, \hat{\hat{t}} \in T^{+}$.

By the cutoff structure established in step 2, $\boldsymbol{x}(\hat{t},\cdot) = (0,\dots,0, x(\hat{t},\tilde{s}(\hat{t})), 1,\dots,1)$ for all $t \in T^{+}$. Suppose to the contrary that $x(\hat{t},\tilde{s}(\hat{t})) + \Sigma_{s>\tilde{s}(\hat{t})} 1 > x(\hat{\hat{t}},\tilde{s}(\hat{\hat{t}})) + \Sigma_{s>\tilde{s}(\hat{\hat{t}})} 1$ for some $\hat{t}, \hat{\hat{t}} \in T^{+}$. Replacing $\boldsymbol{x}(\hat{\hat{t}},\cdot)$ with $\boldsymbol{x}(\hat{t},\cdot)$ does not generate new incentives to misreport, but it increases the principal's expected value, as it increases the allocation probability for profitable types. If feasibility is hurt—that is, $x(\hat{t},s) + z(\hat{\hat{t}},s) > 1$ for some $s \in S$—decrease $z(\hat{t},s)$ until $x(\hat{t},s) + z(\hat{\hat{t}},s) = 1$. This is also a strict improvement for the principal, as she saves verification costs.

A2.4. Step 4

There is an optimal mechanism in the relaxed problem featuring $\boldsymbol{x}(\hat{t},\cdot) = \boldsymbol{x}(\hat{\hat{t}},\cdot)$ for all $\hat{t}, \hat{\hat{t}} \in T^{+} \cup T^{-}$.

Fix some unprofitable type $t \in T^{-}$. By step 1, we have $\boldsymbol{z}(t,\cdot) = 0$. Optimally, the principal wants to choose the lowest possible allocation probability for the unprofitable types. However, she needs to grant him at least the same interim allocation probability that he could achieve by misreporting to be a profitable type $\hat{t} \in T^{+}$ (by steps 2 and 3, we know that $\boldsymbol{x}(\hat{t},\cdot)$ is the same for all $\hat{t} \in T^{+}$). As the signal realization has no effect on the allocation value, the principal is indifferent between any allocation vectors $\mathbf{x}(t,\cdot)$ that induce the same interim allocation probability, $\mathbb{E}[x(t,s)|t]$, that is, $\mathbb{E}[v(t)x(t,s)|t] = v(t)\mathbb{E}[x(t,s)|t]$. Therefore, she can grant the unprofitable types just the same allocation probability they would face if they misreported a profitable type: $\boldsymbol{x}(t,\cdot) = \boldsymbol{x}(\hat{t},\cdot)$.

This step concludes the proof. If the nonverified allocation probability is independent of the type report at each signal, the agent has no incentive to misreport even if he knows the signal. QED

A3. Proof of Proposition 2

A3.1. Step 1

For any $s \in S$, the optimal $(\boldsymbol{x}(\cdot,s), \boldsymbol{z}(\cdot,s))$ can be determined separately, as all constraints involve only allocation and verification probabilities for the same signal realization. The principal's optimal expected value is the weighted sum of the values of the subproblems

$$\max_{(\boldsymbol{x}(\cdot,s),\boldsymbol{z}(\cdot,s))\geq 0} \mathbb{E}_t[x(t,s)v(t) + z(t,s)(v(t)-c)|s],$$

$$\text{subject to } \forall\ t,\hat{t} \in T:\ (\text{EPIC}(s)_{t,\hat{t}}) \quad \text{and} \qquad (\text{LP(s)})$$

$$\forall\ t \in T:\ x(t,s) + z(t,s) \leq 1.$$

A3.2. Step 2

For any $s \in S$ and for all $t,\hat{t} \in T$:, $x(t,s) = x(\hat{t},s)$; that is, the allocation probability $\boldsymbol{x}(\cdot,s)$ has to be constant in the report.

Suppose to the contrary that there were reports t and $\hat{t}$ with $x(\hat{t}, s) > x(t, s)$. EPIC implies that for all $\tilde{t} \in T$, we have $x(\tilde{t}, s) + z(\tilde{t}, s) \geq x(\hat{t}, s) > x(t, s)$. Hence, there cannot be a type with a binding incentive constraint regarding the report t. This, in turn, implies that optimally, $z(t, s) = 0$. If it were positive, $z(t, s)$ could be lowered and $x(t, s)$ could be increased, at least until the strict inequality above binds. This leaves the allocation probabilities unchanged but lowers verification costs. The incentive constraints of type t now take the form $x(t, s) + 0 \geq x(\tilde{t}, s)$ for all reports $\tilde{t}$ and, in particular, for report $\hat{t}$, contradicting the above hypothesis. Hence, we must have that for all $t, \hat{t}$: $x(t, s) = x(\hat{t}, s) \equiv \chi_s$.

A3.3. Step 3

With constant $\boldsymbol{x}(\cdot, s)$, all incentive constraints are automatically fulfilled, as the unverified allocation probability is the same for any possible report. The principal's problem reads as follows:

$$\max_{(\chi_s, \boldsymbol{z}(\cdot, s)) \geq 0} \sum_{t \in T} f(t, s)[\chi_s v(t) + z(t, s)(v(t) - c)], \qquad \text{subject to } \forall\, t \in T: \chi_s + z(t, s) \leq 1. \tag{LP(s)}$$

In this simplified program, $z(t, s)$ will be set as high as possible—that is, to $1 - \chi_s$—if $(v(t, s) - c)$ is positive and to 0 otherwise, yielding the following:

$$\max_{\chi_s \in [0,1]} \chi_s \cdot \sum_{t \in T} f(t, s) v(t) + \sum_{t \in T} f(t, s)(1 - \chi_s)(v(t) - c)^+. \tag{LP(s)}$$

A3.4. Step 4

Expressed in terms of conditional expectations, the problem is linear in χ_s:

$$\max_{\chi_s \in [0,1]} \chi_s \cdot \mathbb{E}_t[v(t)|s] + (1 - \chi_s) \cdot \mathbb{E}_t[(v(t) - c)^+|s]. \tag{LP(s)}$$

Generically, the optimal value of χ_s is either 0 or 1, depending on which of the expectations is larger. The optimality of the cutoff mechanism follows, as $\mathbb{E}_t[v(t)|s] - \mathbb{E}_t[(v(t) - c)^+|s]$ is increasing in s. QED

A4. Proof of Proposition 3

A4.1. Step 1

Similar to the case with perfect verification, we can solve for the optimal transparent mechanism for each signal realization s separately where the incentive constraint now reads

$$x(t, s) + z(t, s) \geq x(\hat{t}, s) + \varepsilon z(\hat{t}, s). \tag{$\text{EPIC}(s)^{\varepsilon}_{t,\hat{t}}$}$$

A4.2. Step 2

For any $s \in S$ and for all $t, t' \in T: x(t,s) + \varepsilon z(t,s) = x(t',s) + \varepsilon z(t',s)$, that is, the expected allocation probability from misreporting is constant in the (mis)report.

Suppose to the contrary that there are types t and t' with $x(t,s) + \varepsilon z(t,s) > x(t',s) + \varepsilon z(t',s)$. Then, no other type can have a binding incentive constraint regarding the misreport t'. It follows that optimally $z(t',s) = 0$. Otherwise, the principal could benefit from decreasing $z(t',s)$ and increasing $x(t',s)$ at the same rate until $z(t',s) = 0$ or $x(t',s) + \varepsilon z(t',s) = x(t,s) + \varepsilon z(t,s)$. However, if $z(t',s) = 0$ and still $x(t,s) + \varepsilon z(t,s) > x(t',s) + \varepsilon z(t',s) = x(t',s) + \varepsilon 0$, then also $x(t,s) + \varepsilon z(t,s) > x(t',s) + z(t',s)$, so that type t' would have a strict incentive to misreport to type t.

A4.3. Step 3

Given that for all t we have $x(t,s) + \varepsilon z(t,s) = \chi_s$, all IC constraints are satisfied. Thus, for any fixed value of $\chi_s \in [0,1]$, the principal chooses the optimal level allocation and verification probabilities for each t, subject to the feasibility constraint $x(t,s) + z(t,s) \leq 1$ and to $x(t,s) + \varepsilon z(t,s) = \chi_s$. That is, the principal chooses $x(t,s) \in [(\chi_s - \varepsilon)^+/(1-\varepsilon), \chi_s]$ and sets $z(t,s) = (\chi_s - x(t,s))/\varepsilon$.[36] Now, take any type t and consider the optimal distribution of χ_s over $x(t,s)$ and $z(t,s)$. The principal's allocation value from type t as a function of $x(t,s)$ is

$$\max_{x \in [(\chi-\varepsilon)^+/(1-\varepsilon), \chi_s]} v(t) \cdot x + (v(t) - c) \cdot \frac{\chi_s - x}{\varepsilon}, \tag{A1}$$

where the last fraction is $z(t,s)$, given that $x(t,s) = x$ and $x + \varepsilon z(t,s) = \chi_s$. It is easy to see that it is optimal for the principal to set

$$x(t,s) = \begin{cases} \chi_s & \text{if } v(t) \leq \dfrac{c}{1-\varepsilon}, \\[2ex] \dfrac{(\chi-\varepsilon)^+}{1-\varepsilon} & \text{if } v(t) > \dfrac{c}{1-\varepsilon}. \end{cases} \tag{A2}$$

In words, there is a threshold type $\bar{t}^{\varepsilon} = \min\{t | v(t) - [c/(1-\varepsilon)] > 0\}$ such that all types below the threshold are allocated without verification only, while types above the threshold are allocated after verification with maximal possible probability, given χ_s, and allocated without verification only when $\chi_s > \varepsilon$.

A4.4. Step 4

Given the optimal choice of $x(t,s)$ and $z(t,s)$ in the previous step, the principal's optimal choice of χ_s solves

$$\max_{\chi \in [0,1]} \chi \cdot \mathbb{E}_t[v(t)\mathbf{1}_{\{t<\bar{t}\}}|s] \tag{A3}$$

$$+ \max\left\{\frac{\chi-\varepsilon}{1-\varepsilon}, 0\right\} \cdot \mathbb{E}_t[v(t)\mathbf{1}_{\{t\geq\bar{t}\}}|s] + \min\left\{\frac{1-\chi}{1-\varepsilon}, \frac{\chi}{\varepsilon}\right\} \cdot \mathbb{E}_t\left[(v(t)-c)\mathbf{1}_{\{t\geq\bar{t}\}}\middle| s\right]. \tag{A4}$$

[36] The lower bound for x, $(\chi-\varepsilon)^+/(1-\varepsilon)$, is required, since whenever $\chi_s > \varepsilon$, at least part of the total expected allocation from misreporting has to come from x rather than z.

The objective is piecewise linear in χ with a kink at $\chi = \varepsilon$, and it is optimal to choose χ_s such that

$$\chi_s = \begin{cases} 1 \text{ if } \mathbb{E}_t[v(t)|s] > \mathbb{E}_t\left[\left(v(t) - \frac{c}{1-\varepsilon}\right)^+ \,\middle|\, s\right], \\ \varepsilon \text{ if } \mathbb{E}_t[v(t)|s] \in \left(-\frac{1-\varepsilon}{\varepsilon}\mathbb{E}_t\left[\left(v(t) - \frac{c}{1-\varepsilon}\right)^+ \,\middle|\, s\right], \mathbb{E}_t\left[\left(v(t) - \frac{c}{1-\varepsilon}\right)^+ \,\middle|\, s\right]\right), \\ 0 \text{ if } \mathbb{E}_t[v(t)|s] \le -\mathbb{E}_t\left[\left(v(t) - \frac{c}{1-\varepsilon}\right)^+ \,\middle|\, s\right]. \end{cases} \tag{A5}$$

The three cases above correspond to the signal ranges defined by the cutoffs $\underline{s}^\varepsilon$ and $\bar{s}^\varepsilon$ in proposition 3. QED

A5. Proof of Proposition 4

Fix an environment with noisy verification technology α such that $\varepsilon = \max_{\hat{t},t\in T:\hat{t}\neq t}\alpha(\hat{t}|t)$ constitutes the largest verification error across all type-misreport combinations. To establish the bound on the principal's gain from secrecy given in the proposition, we give a lower bound on the principal's value from the optimal transparent—that is, EPIC—mechanism and an upper bound on the principal's value from the optimal mechanism in the larger class of BIC mechanisms.

A5.1. Step 1

Denote by $V_\varepsilon^{\text{trans}}$ the maximal payoff the principal can achieve from a transparent mechanism. This value is determined by the linear program (LP) with the set of IC constraints replaced by

$$x(t, s) + z(t, s) \ge x(\hat{t}, s) + \alpha(\hat{t}|t)z(\hat{t}, s), \text{ for all } t, \hat{t}, s. \tag{A6}$$

Let $(\boldsymbol{x}, \boldsymbol{z})$ be the optimal cutoff-with-appeal mechanism in the case of perfect verification in theorem 2, and denote the principal's value from that mechanism by V_0^{trans}, where the subscript 0 indicates that the verification error is 0 in that case of perfect verification. We apply two slight modifications to the original mechanism $(\boldsymbol{x}, \boldsymbol{z})$. First, define a modified mechanism $(\boldsymbol{x}', \boldsymbol{z}')$ such that $x'(t, s) = \varepsilon$ for all (t, s), with $t < \bar{t}$ and $s < \bar{s}$, and $x'(t, s) = x(t, s)$, $z'(t, s) = z(t, s)$ for all remaining (t, s). That is, the modified mechanism introduces (unverified) allocation probability ε at all those type-signal combinations that lead to zero allocation probability in the original cutoff-with-appeal mechanism under perfect verification. Under mechanism $(\boldsymbol{x}', \boldsymbol{z}')$, the principal's conditional expected value after any signal realization $s < \bar{s}$ is equal to

$$\mathbb{E}_t[(v(t) - c)^+|s] + \varepsilon\mathbb{E}_t[v(t)\mathbf{1}_{\{t<\bar{t}\}}|s], \tag{A7}$$

where the second term captures the change due to the introduced modification. Note that the definition of $\bar{s}$ implies that the second expectation is negative for any $s < \bar{s}$. In particular, it may be that the entire expected value conditional on s is negative for low enough signal realizations.

After such signals, the principal is better off not allocating to the agent at all and never verifying independent of the agent's report. Formally, construct the modification $(\tilde{\mathbf{x}}, \tilde{\mathbf{z}})$ by setting $\tilde{x}(t, s) = \tilde{z}(t, s) = 0$ for all t and all s such that value (A7) is negative, maintaining $\tilde{x}(t, s) = x'(t, s)$ and $\tilde{z}(t, s) = z'(t, s)$ for all other (t, s).

It is easily verified that the modified mechanism $(\tilde{\mathbf{x}}, \tilde{\mathbf{z}})$ satisfies the IC constraints in (A6), given that $\alpha(\hat{t}|t) \leq \varepsilon$ for all combinations $\hat{t} \neq t$. Therefore, the modification is feasible in the principal's problem with verification error, and its value constitutes a lower bound for the principal's value from any incentive-compatible transparent mechanism under this noisy verification technology. Denote the principal's expected value from the modified mechanism by $\tilde{V}_\varepsilon$, for which we just showed that $V_\varepsilon^{\text{trans}} \geq \tilde{V}_\varepsilon$

A5.2. Step 2

Denote by $V_\varepsilon^{\text{secr}}$ the maximal payoff the principal can achieve from any (i.e., potentially nontransparent) mechanism. This value is given by the linear program (LP) with the set of IC constraints replaced by

$$\mathbb{E}_s[x(t, s) + z(t, s) - x(\hat{t}, s) - \alpha(\hat{t}|t)z(\hat{t}, s)|t] \geq 0, \text{ for all } t, \hat{t}. \tag{A8}$$

The above IC constraints are clearly more restrictive than the original $(\text{BIC}_{t,\hat{t}})$ constraints in (LP). As the objective of the program is unchanged, it follows that the principal's value from the optimal mechanism with verification error must be bounded by the value from the optimal mechanism with perfect verification. Intuitively, the principal with perfect verification technology could always commit to "simulate" any verification error internally by changing her reaction to verification outcomes with appropriate probabilities. Thus, we must have $V_\varepsilon^{\text{secr}} \leq V_0^{\text{secr}}$, where the latter denotes the principal's value from the optimal cutoff-with-appeal mechanism in the case of perfect verification.

A5.3. Step 3

The gain from secrecy is defined as $V_\varepsilon^{\text{secr}} - V_\varepsilon^{\text{trans}}$. By steps 2 and 1 above, we have $V_\varepsilon^{\text{secr}} - V_\varepsilon^{\text{trans}} \leq V_0^{\text{secr}} - \tilde{V}_\varepsilon$. Finally, our transparency result proposition 1 tells us that $V_0^{\text{secr}} = V_0^{\text{trans}}$. The claim in proposition 4 follows by verifying that

$$V_0^{\text{trans}} - \tilde{V}_\varepsilon = \varepsilon \cdot \sum_{s<\bar{s}} \mathbb{P}[s] \min\left\{\frac{1}{\varepsilon}\mathbb{E}[(v(t) - c)^+|s]; \mathbb{E}[-v(t)\mathbf{1}_{t<\hat{t}}|s]\right\}.$$

QED

A6. Proof of Proposition 5

We present the following relaxation of the principal's problem under coarse verification and show that it is solved by a transparent mechanism that is also feasible in the original program. Divide the partition $\mathcal{T}$ into two or three subsets. The set of partition elements for which all types are profitable is denoted

$$\mathcal{T}^+ \equiv \{\tau \in \mathcal{T} | v(t) > 0 \ \ \forall \ t \in \tau\}.$$

The set of partition elements for which all types are unprofitable is denoted

$$\mathcal{T}^- \equiv \{\tau \in \mathcal{T} | v(t) \leq 0 \ \ \forall \ t \in \tau\}.$$

Finally, there may be one mixed partition element, $\tau^{\text{mix}} \in \mathcal{T}/(\mathcal{T}^+ \cup \mathcal{T}^-)$, that contains both profitable and unprofitable types. Note that we can restrict attention to those cases in which at most one of the three sets above may be empty. If two are empty, the principal-optimal allocation is either constant (all types have positive or all types have negative value) or no verification is possible (all types are contained in τ^{mix}). We slightly abuse notation below and write $t \in \mathcal{T}^+$ to express that $t \in \tau$ for some $\tau \in \mathcal{T}^+$ and likewise for $\mathcal{T}^-$.

The relaxed problem includes only upward incentive constraints, but not all of them. We ignore all constraints that prevent misreports to types in $\mathcal{T}^-$ from types in other partition elements. Hence, it reads as follows:

$$\begin{aligned}
&\max_{x,z\geq 0} \sum_{t\in T}\sum_{s\in S} f(t,s)[x(t,s)v(t) + z(t,s)(v(t)-c)],\\
&\quad\text{subject to } \forall (t,\hat{t}) \in \mathcal{T} \times \mathcal{T}^+ \cup \{\tau^{\text{mix}}\} \text{ with } t < \hat{t} : (\text{BIC}_{t,\hat{t}}),\\
&\qquad\qquad \forall (t,\hat{t}) \in T \times T \text{ with } t < \hat{t} \text{ and } \hat{t} \in \tau(t) : (\text{CBIC}_{t,\hat{t}}),\\
&\qquad\qquad \forall (t,s) \in T \times S : x(t,s) + z(t,s) \leq 1.
\end{aligned} \tag{CLP.r}$$

A6.1. Step 1

In any solution to the relaxed problem, $z(t,s) = 0$ for all $t \in \mathcal{T}^-$ and all $s \in S$.

Suppose otherwise that for some $t \in \mathcal{T}^-$ and some $s \in S$ we have $z(t,s) > 0$. Shifting probability mass from $z(t, s)$ to $x(t, s)$ such that the overall allocation probability stays constant,

$$0 < dx(t,s) = -dz(t,s),$$

saves the principal verification costs and does not distort the incentives. Type t's incentives to misreport remain the same, and the only incentive constraints to misreport type $t \in \mathcal{T}^-$ included in the relaxed problem are the intrainterval (CBIC) constraints for which only the sum $x(t,s) + z(t,s)$ matters.

A6.2. Step 2

There is an optimal mechanism in the relaxed problem featuring a cutoff for $\boldsymbol{x}(\hat{t},\cdot)$:

$$\forall\ \hat{t} \in \mathcal{T}^+ \cup \{\tau^{\text{mix}}\}\ \exists\ \tilde{s}(\hat{t}) \in S : x(\hat{t},s) \begin{cases} = 0 & \text{if } s < \tilde{s}(\hat{t}),\\ \in [0,1) & \text{if } s > \tilde{s}(\hat{t}),\\ = 1 & \text{if } s = \tilde{s}(\hat{t}). \end{cases}$$

Take a feasible mechanism in the relaxed problem and suppose that for some $\hat{t} \in \mathcal{T}^+ \cup \{\tau^{\text{mix}}\}$, $\exists\ s < s'$ such that $x(\hat{t},s) > 0$, $x(\hat{t},s') < 1$. Modify this mechanism at two points, shifting the allocation probability from $x(\hat{t},s)$ to $x(\hat{t},s')$, that is, $dx(\hat{t},s) < 0$ and $dx(\hat{t},s') > 0$. Choose these shifts in a proportion such that for $\hat{t}$, the expected allocation probability remains unchanged:

$$0 = f(\hat{t}, s)dx(\hat{t}, s) + f(\hat{t}, s')dx(\hat{t}, s') \Leftrightarrow dx(\hat{t}, s) = -\frac{f(\hat{t}, s')}{f(\hat{t}, s)}dx(\hat{t}, s').$$

For all types $t < \hat{t}$, this reduces the incentive to report $\hat{t}$ because

$$\begin{aligned} d(\mathrm{BIC}_{t,\hat{t}}) &= -f(t, s)dx(\hat{t}, s) - f(t, s')dx(\hat{t}, s') \\ &= f(t, s)\left[\frac{f(\hat{t}, s')}{f(\hat{t}, s)} - \frac{f(t, s')}{f(t, s)}\right]dx(\hat{t}, s') > 0 \end{aligned}$$

by the MLRP.

The proposed shift is feasible if in the original mechanism, $x(\hat{t}, s') + z(\hat{t}, s') < 1$. In the case that $x(\hat{t}, s') + z(\hat{t}, s') = 1$, it can still be implemented by shifting in addition mass from $z(\hat{t}, s')$ to $z(\hat{t}, s)$ to ensure that $x(\hat{t}, s') + z(\hat{t}, s')$ and $x(\hat{t}, s) + z(\hat{t}, s)$ remain constant:

$$dx(\hat{t}, s') + dz(\hat{t}, s') = 0 \quad \text{and} \quad dx(\hat{t}, s) + dz(\hat{t}, s) = 0.$$

This implies that $dz(\hat{t}, s') < 0$ and $dz(\hat{t}, s) > 0$. This is feasible, as $x(\hat{t}, s') < 1$ and $x(\hat{t}, s') + z(\hat{t}, s') = 1$ imply that $z(\hat{t}, s') > 0$. As $x(\hat{t}, s) > 0$, we must further have $z(\hat{t}, s) < 1$ by feasibility. To maintain the total allocation probabilities constant, the above changes in x are compensated for by the following changes in z:

$$dz(\hat{t}, s) = \frac{f(\hat{t}, s')}{f(\hat{t}, s)}(-dz(\hat{t}, s')).$$

The incentives for any lower type from a different partition element to misreport type $\hat{t}$ are weakened in the same way as above because $z(\hat{t}, s)$ and $z(\hat{t}, s')$ do not play a role in the constraints that prevent misreport $\hat{t}$. For the types in the same partition element, the total allocation $x + z$ has not changed, so their incentive to report $\hat{t}$ remains unchanged.

A6.3. Step 3

There is an optimal mechanism in the relaxed problem featuring a cutoff in $z(\hat{t}, \cdot)$:

$$\forall\, \hat{t} \in \mathcal{T}^+ \cup \{\tau^{\text{mix}}\}\ \exists\ \underline{s}(\hat{t}) \in S \text{ with } \underline{s}(\hat{t}) \le \tilde{s}(\hat{t}),$$

$$\text{subject to } z(\hat{t}, s)\begin{cases} = 0 & \text{if } s \notin [\underline{s}(\hat{t}), \tilde{s}(\hat{t})], \\ \in [0, 1) & \text{if } s = \underline{s}(\hat{t}), \\ = 1 & \text{if } \underline{s}(\hat{t}) < s < \tilde{s}(\hat{t}), \\ = 1 - x(\hat{t}, s) & \text{if } \underline{s}(\hat{t}) < s = \tilde{s}(\hat{t}). \end{cases}$$

Given the cutoff in x from the previous step, apply the same shifts from lower to higher signals to verification probability z. For any report $\hat{t}$, the incentive constraints for true types not in $\tau(\hat{t})$ are unaffected. The upward constraints for any type $t \in \tau(\hat{t})$ with $t < \hat{t}$ are relaxed. The principal's payoff is unchanged.

A6.4. Step 4

The $\boldsymbol{x}$ cutoff is the same for all profitable types, that is, $x(t, s) = x(t', s)$ for all $t, t' \in \mathcal{T}^+$ and all $s \in S$. Further, $x(t, s) \geq x(\tilde{t}, s)$ for all $(t, \tilde{t}) \in \mathcal{T}^+ \times \tau^{\text{mix}}$.

Take the report $\hat{t}$ among profitable types with the largest x allocation (with the signal cutoff farthest to the left); that is, take $\hat{t} = \arg\min_{t \in \mathcal{T}^+}(\tilde{s}(t))$. If the set is not a singleton, choose any among the ones for which $x(\hat{t}, \tilde{s}(\hat{t}))$ is maximal. Now set $x(t, s) = x(\hat{t}, s)$ for all $t \in \mathcal{T}^+$ and all $s \in S$. If this change results in $x(t, s) + z(t, s) > 1$, reduce verification to $z(t, s) = 1 - x(\hat{t}, s)$. This change is profitable for the principal, as it increases allocation probability to profitable types and decreases verification probability. The resulting mechanism is also incentive compatible in the relaxed problem. For any report $t \in \mathcal{T}^+$, any type $\tilde{t} \notin \tau(t)$ gets the same outcome from reporting t he could already get by reporting $\hat{t}$ before the modification. For lower types $\tilde{t} < t$ in the same partition element $\tau(t)$, the cutoff structure implies that for all s, either $x(t, s) + z(t, s)$ remains unchanged or $z(\tilde{t}, s) = z(t, s) = 0$ and $x(\tilde{t}, s) = x(t, s)$, so that incentive compatibility is preserved by this modification.

Finally, note that any optimal mechanism in the relaxed problem must feature $x(t, s) \geq x(\tilde{t}, s)$ for all $t \in \mathcal{T}^+$ and $\tilde{t} \in \tau^{\text{mix}}$. If this were not the case, we could increase x on $\mathcal{T}^+$ without violating any IC constraints and thereby either increase the allocation probability for strictly positive types or decrease the verification probability.

A6.5. Step 5

For all $t, t' \in \tau^{\text{mix}}$, we have $x(t, s) = x(t', s)$ and $z(t, s) = z(t', s)$ for all s.

The cutoff structure from steps 2 and 3, together with the upward incentive constraints within τ^{mix}, imply that for any type pair $t < t'$ within τ^{mix}, we must have $x(t, s) + z(t, s) \geq x(t', s) + z(t', s)$ for all s.

We first show that any solution in the relaxed problem must have $x(t', s) \geq x(t, s)$. Suppose that this was not the case. Then there must be a pair of consecutive types $t' < t''$ with $x(t'', s) < x(t', s)$. Let t', t'' be the lowest such pair. Then it must be that $z(t'', s) = 0$ for all s, because no type from a lower partition element can have a binding IC constraint to report t'' without having a strict incentive to misreport t'. This, in turn, implies that for all types $t \in \tau^{\text{mix}}$ with $t > t''$, we have $z(t, s) = 0$ for all s. That is because $x(t, s) + z(t, s) \leq x(t'', s) + z(t'', s) = x(t'', s) < x(t', s)$, so no binding constraint from another partition element could justify the verification cost. Now by $t' < t''$, either $v(t'') > 0$ or $v(t') \leq 0$. In the first case, if $v(t'') > 0$, then optimality requires that $x(t'', s) = x(t', s)$, contradicting $x(t'', s) < x(t', s)$. In the second case, if $v(t') \leq 0$, then $v(\tilde{t}) - c < 0$ for all $\tilde{t} \leq t'$. In that case, optimality requires that $z(\tilde{t}, s) = 0$ for all $\tilde{t} \leq t'$ and all s. By incentive compatibility and given that $x(t, s)$ is monotone in t below t', it must hold that $z(t', s) = 0$ and $x(\tilde{t}, s) = x(t', s)$ for all $\tilde{t} \leq t'$ and all s. Again, incentive compatibility then specifies that $x(t', s) = x(t'', s)$ (so that all x are equal on τ^{mix}), again contradicting $x(t'', s) < x(t', s)$. Hence, within τ^{mix}, $x(t, s)$ must be monotone increasing in t for all s. It follows immediately that $x(t, s)$ must be constant in t for all s. If this were not the case, there would be $t < t'$ with $x(t, s) < x(t', s)$ for some s. For t not to misreport t', this requires that $z(t, s) > 0$. This cannot be optimal, given that replacing $z(t, s)$ with $x(t, s)$ would be incentive compatible and save verification costs.

Second, we show that $z(t, s) = z(t', s)$ for all t and t' in τ^{mix} and all s. Given that x is constant on τ^{mix}, the upward incentive constraints within τ^{mix} imply that the $z(t, s)$ must be weakly decreasing in t within τ^{mix} for all s. Suppose that $z(t, s)$ is strictly decreasing in t for some s; that is, for $t < t'$, we have $z(t, s) > z(t', s)$. Since $t < t'$, either $v(t) - c \leq 0$ or $v(t') - c > 0$. In the first case, it is optimal to decrease $z(t, s)$; in the latter case, it is optimal to increase $z(t', s)$. Both changes preserve incentive compatibility as long as we maintain $z(t, s) \geq z(t', s)$.

A6.6. Step 6

The $\boldsymbol{x}$ cutoff is the same for $\mathcal{T}^+$ and τ^{mix}.

For all s, the $x(t, s)$ are constant in t both on $\mathcal{T}^+$ and on τ^{mix}. Suppose that the two are different. By step 4, then there must be some s such that for $(t, t') \in \tau^{\text{mix}} \times \mathcal{T}^+$: $x(t, s) < x(t', s)$. Then, incentive compatibility for type t in τ^{mix} requires that $z(t, s) > 0$. This cannot be optimal, given that replacing $z(t, s)$ with $x(t, s)$ would be incentive compatible and save verification costs.

A6.7. Step 7: Cutoff in $\boldsymbol{x}$ for $\mathcal{T}^-$

Finally, since $v(t) < 0$ for all types in $\mathcal{T}^-$, it is easy to see that their allocation x must be chosen as low as permitted by incentive compatibility. This results in $x(t, s) = x(t', s)$ for all pairs $(t, t') \in T \times T$ and all s.

With the last step, we arrive at an optimal mechanism for which the nonverified allocation probability x is constant for all types and the verified allocation probability z is constant within all partition elements. Any such mechanism clearly satisfies the constraints that were ignored in the relaxed problem, so that it solves the original problem. Furthermore, any such mechanism is transparently implementable; that is, it satisfies the pointwise constraints $(\text{EPIC}(s)_{t,\hat{t}})$ for all $(t, \hat{t})$ as well as the pointwise constraints analogous to $(\text{CBIC}_{t,\hat{t}})$ for all pairs $(t, \hat{t})$ with $\hat{t} \in \tau(t)$.

To derive the optimal mechanism given in proposition 5, we can establish the principal-optimal mechanism under coarse verification for each signal s separately. The steps are analogous to those in the proof of proposition 2. Finally, we verify that the cutoffs s_c^* and $\bar{s}_c$ are well defined. For s_c^*, it follows immediately from the MLRP that $\mathbb{E}_t[(v(t) - c)\mathbf{1}_{\{t \in \tau(\bar{t})\}} | s]$ crosses 0 at most once. To verify this, we can divide the expression by $\mathbb{P}[t \in \tau(\bar{t}) | s] > 0$ and confirm that the conditional expectation $\mathbb{E}_t[(v(t) - c) | t \in \tau(\bar{t}); s]$ is increasing in s. For $\bar{s}_c$, we can write the expression in the definition,

$$\mathbb{E}_t[v(t) | s] - \mathbb{E}_t[(v(t) - c)\mathbf{1}_{\{t > \max(\tau(\bar{t}))\}} | s] - (\mathbb{E}_t[(v(t) - c)\mathbf{1}_{\{t \in \tau(\bar{t})\}} | s])^+$$

as

$$\min\{\mathbb{E}_t[v(t) | s] - \mathbb{E}_t[(v(t) - c)\mathbf{1}_{\{t > \max(\tau(\bar{t}))\}} | s]; \mathbb{E}_t[v(t) | s] - \mathbb{E}_t[(v(t) - c)\mathbf{1}_{\{t \geq \min(\tau(\bar{t}))\}} | s]\}.$$

This expression crosses 0 at most once, because both elements in the minimum operator cross 0 at most once and from below. We show this for the case of the second expression; the steps for the first expression follow analogously. We can rewrite

$$\mathbb{E}_t[v(t)|s] - \mathbb{E}_t[(v(t)-c)\mathbf{1}_{\{t\geq\min(\tau(\bar{t}))\}}|s] = \mathbb{E}_t[v(t)\mathbf{1}_{\{t<\min(\tau(\bar{t}))\}}|s] + c\cdot\mathbb{P}[t\geq\min(\tau(\bar{t}))|s].$$

Dividing this expression by $1-\mathbb{P}[t\geq\min(\tau(\bar{t}))|s] = \mathbb{P}[t\geq\min(\tau(\bar{t}))|s] > 0$ gives

$$\mathbb{E}_t[v(t)|t<\min(\tau(\bar{t}));s] + c\cdot\frac{\mathbb{P}[t\geq\min(\tau(\bar{t}))|s]}{1-\mathbb{P}[t\geq\min(\tau(\bar{t}))|s]}.$$

Both terms are increasing in s, as a result of the MLRP, so that it can cross 0 at most once and from below. QED

Appendix B

Revelation Principle with Noisy Verification

We consider the following model of noisy verification. For any pair of type t and type report $\hat{t}$, we specify a probability, $\alpha(\hat{t}|t)$, of type t passing as type $\hat{t}$. If the agent reports truthfully, he passes the verification process for sure, $\alpha(t|t) = 1$. We study direct mechanisms that specify, for any type report t and signal realization s,

- a probability of verifying the report: $e(\hat{t}, s)$;
- an allocation decision when the report is not verified: $a(\hat{t}, s, \varnothing)$;
- an allocation when the type passes as the reported type: $a(\hat{t}, s, 1)$;
- and an allocation when the reported type fails the verification: $a(\hat{t}, s, 0)$.

Ball and Kattwinkel (2019) show how this reduced-form model can be microfounded by a setting in which—in addition to the standard form of cheap-talk communication—there exists a set of pass/fail tests (with type-dependent passing probabilities) from which the principal can choose one and conduct it on the agent.[37] They show that in this setting it is furthermore without loss to restrict to direct mechanisms that induce truthful reporting. With that, our focus on the direct mechanisms introduced above is without loss of generality.

Analogously to the case with perfect verification, we can use optimality arguments to restrict the set of relevant mechanisms further.

Lemma 3 (maximal punishment and minimal verification). It is without loss of optimality for the principal to restrict the class of direct mechanisms to fulfill the following properties:

1. maximal punishment: $a(t, s, 0) = 0$; and
2. minimal verification: if $e(t, s) > 0$, then $a(t, s, 1) = 1$.

Proof. First, since $\alpha(t|t) = 1$, misreporting is an off-path event, and therefore the allocation probability after a misreport is detected, $a(t, s, 0)$, does not affect the

[37] They show that the microfoundation inherently puts a structure on the passing rate: $\alpha(t_3|t_1) \geq \alpha(t_2|t_1)\alpha(t_3|t_2)$.

principal's payoff. For the agent, setting the allocation after a misreport as low as possible decreases incentives to misreport while maintaining the payoff of any truth-telling type unchanged.

Second, suppose that $e(t, s) > 0$ and $a(t, s, 1) < 1$. One could now lower the probability of verification $e(t, s) \downarrow$ while increasing the probability of allocation after confirming the report as true $a(s, t, 1) \uparrow$, such that $e(t, s)a(t, s, 1)$ remains constant. Note that for any misreporting type $\tilde{t}$, $e(t, s)a(t, s, 1)\alpha(t|\tilde{t})$ also remains constant by this operation.

Lowering the verification probability would, therefore, increase the incentives to misreport and the overall allocation probability only after report t and signal s, if there was allocation with positive probability conditional on no verification, that is, $a(t, s, \varnothing) > 0$. However, in this case, this allocation could be lowered $a(t, s, \varnothing) \downarrow$, such that $(1 - e(t, s))a(t, s, \varnothing)$ remains constant. Then, the incentives to misreport and the overall allocation probability would remain constant. As these procedures would save verification costs while keeping all unconditional allocation probabilities constant, the fact that an optimal mechanism features nonmaximal reward can be ruled out. QED

References

Akbarpour, Mohammad, and Shengwu Li. 2020. "Credible Auctions: A Trilemma." *Econometrica* 88 (2): 425–67.

Ball, Ian, and Deniz Kattwinkel. 2019. "Probabilistic Verification in Mechanism Design." In *EC '19: Proceedings of the 2019 ACM Conference on Economics and Computation*, 389–90. New York: Assoc. Computing Machinery. https://doi.org/10.1145/3328526.3329657.

Ben-Porath, Elchanan, Eddie Dekel, and Barton L. Lipman. 2014. "Optimal Allocation with Costly Verification." *A.E.R.* 104 (12): 3779–813.

———. 2019. "Mechanisms with Evidence: Commitment and Robustness." *Econometrica* 87 (2): 529–66.

Beshkar, Mostafa, and Eric W. Bond. 2017. "Cap and Escape in Trade Agreements." *American Econ. J. Microeconomics* 9 (4): 171–202.

Bhargava, Mohit, Dipjyoti Majumdar, and Arunava Sen. 2015. "Incentive-Compatible Voting Rules with Positively Correlated Beliefs." *Theoretical Econ.* 10 (3): 867–85.

Börgers, Tilman, and Peter Postl. 2009. "Efficient Compromising." *J. Econ. Theory* 144 (5): 2057–76.

Bull, Jesse, and Joel Watson. 2004. "Evidence Disclosure and Verifiability." *J. Econ. Theory* 118 (1): 1–31.

Bushman, Robert M., Joseph D. Piotroski, and Abbie J. Smith. 2004. "What Determines Corporate Transparency?" *J. Accounting Res.* 42 (2): 207–52.

Cella, Michela. 2008. "Informed Principal with Correlation." *Games and Econ. Behavior* 64 (2): 433–56.

Crémer, Jacques, and Richard P. McLean. 1988. "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions." *Econometrica* 56 (6): 1247–57.

Demougin, Dominique M., and Devon A. Garvie. 1991. "Contractual Design with Correlated Information under Limited Liability." *Rand J. Econ.* 22 (4): 477–89.

Deneckere, Raymond, and Sergei Severinov. 2008. "Mechanism Design with Partial State Verifiability." *Games and Econ. Behavior* 64 (2): 487–513.

Epitropou, Markos, and Rakesh Vohra. 2019. "Optimal On-Line Allocation Rules with Verification." In *Algorithmic Game Theory: 12th International Symposium, SAGT 2019*, edited by Dimitris Fotakis and Evangelos Markakis, 3–17. Cham: Springer. https://doi.org/10.1007/978-3-030-30473-7_1.

Erlanson Albin, and Andreas Kleiner. 2020. "Costly Verification in Collective Decisions." *Theoretical Econ.* 15 (3): 923–54.

Eső, Péter, and Balázs Szentes. 2007. "Optimal Information Disclosure in Auctions and the Handicap Auction." *Rev. Econ. Studies* 74 (3): 705–31.

Franz II. 1803. *Gesetzbuch über Verbrechen und Schwere Polizey-Uebertretungen vom 3. September 1803*, vol. 2, *Von den schweren Polizey-Uebertretungen*, 115–25. Vienna: Johann Georg Edlen von Mösle.

Gale, Douglas, and Martin Hellwig. 1985. "Incentive-Compatible Debt Contracts: The One-Period Problem." *Rev. Econ. Studies* 52 (4): 647–63.

Gershkov, Alex, Jacob K. Goeree, Alexey Kushnir, Benny Moldovanu, and Xianwen Shi. 2013. "On the Equivalence of Bayesian and Dominant Strategy Implementation." *Econometrica* 81 (1): 197–220.

Gibbard, Allan. 1973. "Manipulation of Voting Schemes: A General Result." *Econometrica* 41 (4): 587–601.

Glazer, Jacob, and Ariel Rubinstein. 2004. "On Optimal Rules of Persuasion." *Econometrica* 72 (6): 1715–36.

Goldlücke, Susanne, and Thomas Tröger. 2018. "Assigning an Unpleasant Task without Payment." Working Paper no. 18-02, Dept. Econ., Univ. Mannheim.

Green, Jerry R., and Jean-Jacques Laffont. 1986. "Partially Verifiable Information and Mechanism Design." *Rev. Econ. Studies* 53 (3):447–56.

Halac, Marina, and Pierre Yared. 2020. "Commitment versus Flexibility with Costly Verification." *J.P.E.* 128 (12): 4523–73.

Hart, Sergiu, Ilan Kremer, and Motty Perry. 2017. "Evidence Games: Truth and Commitment." *A.E.R.* 107 (3): 690–713.

Hood, Christopher, and David Heald, editors. 2006. *Transparency: The Key to Better Governance?* Proc. British Acad., vol. 135. Oxford: Oxford Univ. Press (for British Acad.).

Jehiel, Philippe. 2015. "On Transparency in Organizations." *Rev. Econ. Studies* 82 (2): 736–61.

Johnson, Scott, John W. Pratt, and Richard J. Zeckhauser. 1990. "Efficiency despite Mutually Payoff-Relevant Private Information: The Finite Case." *Econometrica* 58 (4): 873–900.

Kartik, Navin, and Olivier Tercieux. 2012. "Implementation with Evidence." *Theoretical Econ.* 7 (2): 323–55.

Kattwinkel, Deniz. 2020. "Allocation with Correlated Information: Too Good to Be True." In *EC '20: Proceedings of the 21st ACM Conference on Economics and Computation*, 109–10. New York: Assoc. Computing Machinery. https://doi.org/10.1145/3391403.3399536.

Kittler, Wolf. 2003. "Heimlichkeit und Schriftlichkeit: Das österreichische Strafprozessrecht in Franz Kafkas Roman *Der Proceß*." *Germanic Review* 78 (3): 194–222.

Koessler, Frédéric, and Eduardo Perez-Richet. 2019. "Evidence Reading Mechanisms." *Social Choice and Welfare* 53 (3): 375–97.

Li, Yunan. 2020. "Mechanism Design with Costly Verification and Limited Punishments." *J. Econ. Theory* 186:105000.

Manelli, Alejandro M., and Daniel R. Vincent. 2010. "Bayesian and Dominant-Strategy Implementation in the Independent Private-Values Model." *Econometrica* 78 (6): 1905–38.

Maskin, Eric, and Jean Tirole. 1990. “The Principal-Agent Relationship with an Informed Principal: The Case of Private Values.” *Econometrica* 58 (2): 379–409.
McAfee, R. Preston, and Philip J. Reny. 1992. “Correlated Information and Mechanism Design.” *Econometrica* 60 (2): 395–421.
Milgrom, Paul R., and Robert J. Weber. 1982. “A Theory of Auctions and Competitive Bidding.” *Econometrica* 50 (5): 1089–122.
Mookherjee, Dilip, and Ivan Png. 1989. “Optimal Auditing, Insurance, and Redistribution.” *Q.J.E.* 104 (2): 399–415.
Myerson, Roger B. 1981. “Optimal Auction Design.” *Math. Operations Res.* 6 (1): 58–73.
———. 1983. “Mechanism Design by an Informed Principal.” *Econometrica* 51 (6): 1767–97.
Mylovanov, Tymofiy, and Andriy Zapechelnyuk. 2017. “Optimal Allocation with Ex Post Verification and Limited Penalties.” *A.E.R.* 107 (9): 2666–94.
Neeman, Zvika. 2004. “The Relevance of Private Information in Mechanism Design.” *J. Econ. Theory* 117 (1): 55–77.
Ortoleva, Pietro, Evgenii Safonov, and Leeat Yariv. 2021. “Who Cares More? Allocation with Diverse Preference Intensities.” Working Paper no. 29208 (September), NBER, Cambridge, MA.
Ottaviani, Marco, and Andrea Prat. 2001. “The Value of Public Information in Monopoly.” *Econometrica* 69 (6): 1673–83.
Prat, Andrea. 2005. “The Wrong Kind of Transparency.” *A.E.R.* 95 (3): 862–77.
Rappoport, Daniel. 2020. “Evidence and Skepticism in Verifiable Disclosure Games.” Working paper. http://dx.doi.org/10.2139/ssrn.2978288.
Riordan, Michael H., and David E. M. Sappington. 1988. “Optimal Contracts with Public Ex Post Information.” *J. Econ. Theory* 45 (1): 189–99.
Satterthwaite, Mark Allen. 1975. “Strategy-Proofness and Arrow’s Conditions: Existence and Correspondence Theorems for Voting Procedures and Social Welfare Functions.” *J. Econ. Theory* 10 (2): 187–217.
Severinov, Sergei. 2008. “An Efficient Solution to the Informed Principal Problem.” *J. Econ. Theory* 141:114–33.
Siegel, Ron, and Bruno Strulovici. 2020. “Judicial Mechanism Design.” Manuscript, Dept. Econ., Pennsylvania State Univ.
Silva, Francisco. 2019. “If We Confess Our Sins.” *Internat. Econ. Rev.* 60 (3): 1389–412.
———. 2021. “Information Transmission in Persuasion Models with Imperfect Verification.” Working paper. http://dx.doi.org/10.2139/ssrn.3728325.
Skreta, Vasiliki. 2011. “On the Informed Seller Problem: Optimal Information Disclosure.” *Rev. Econ. Design* 15 (1): 1–36.
StPO. 1975. “Die österreichische Strafprozeßordnung: §6 StPO Rechtliches Gehör.” www.jusline.at/gesetz/stpo/paragraf/6.
Townsend, Robert M. 1979. “Optimal Contracts and Competitive Markets with Costly State Verification.” *J. Econ. Theory* 21 (2): 265–93.