Discussion Paper Series – CRC TR 224

Discussion Paper No. 011
Project B 03

Full Surplus Extraction in Mechanism Design With Information Disclosure

Daniel Krähmer [1]

April 2018, revised November 2018

[1] University Bonn, Department of Economics, Institute for Microeconomics, Adenauer Allee 24-42, D-53113 Bonn (Germany), kraehmer@hcm.uni-bonn.de

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

Full surplus extraction in mechanism design with information disclosure

Daniel Krähmer*

November, 2018

Abstract

I study mechanism design settings with quasi-linear utility where the principal can provide agents with additional private information about their valuations beyond the private information they hold at the outset. I demonstrate that the principal can design information and a mechanism so as to fully extract the complete information first-best surplus if agents' ex ante information only affects their beliefs about, yet not their valuations. Otherwise, the result holds if each agent's initial private beliefs satisfy a spanning condition.

Keywords: information design, mechanism design, quasi-linear utility, rent extraction
JEL codes: D82, H57

*University Bonn, Department of Economics, Institute for Microeconomics, Adenauer Allee 24-42, D-53113 Bonn (Germany), kraehmer@hcm.uni-bonn.de. I thank Paul Heidhues, Stephan Lauermann, Benny Moldovanu, Stephen Morris, Martin Pollrich, Ludovic Renou, David Rodina, Alex Smolin, Nicolas Schutz, Roland Strausz, Balazs Szentes, and Tymon Tatur for inspiring discussions. Thanks also go to seminar audiences at Berlin, Bern, Bonn, LSE, UCL. I am grateful to the UCL Economics department for their hospitality. Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 is gratefully acknowledged.

1 Introduction

In many mechanism design settings, a principal can disclose additional information to agents prior to implementing an allocation. For example, sellers often offer interested buyers the possibility to try out, test, or inspect their products prior to purchase. Car dealers offer test drives, online stores offer free book previews, music samples, or movie trailers, or grant consumers withdrawal periods to try out an order for a while. Likewise, an auctioneer of an oil well may allow interested bidders to conduct test drills, or procurement agencies provide contractors with additional information about the costs of carrying out a project.

In this paper, I study mechanism design problems with quasi-linear utility where the principal can, next to a mechanism, design and disclose additional information that affects agents' valuations. Similar to Esö and Szentes (2007a,b) and Li and Shi (2017a), I focus on two issues: First, I focus on situations in which the information the designer provides becomes an agent's private information. For example, in a selling context, trial periods or product descriptions enable buyers primarily to better ascertain whether the product fits their tastes (rather than to verify its objective quality). While a seller may control how much a buyer can possibly learn, for example by setting the time a buyer is allowed to try the product or the richness of the product description, how exactly the information influences a buyer's valuation is not known to the seller but rather becomes the buyer's private information. Second, in addition to the information provided by the principal, agents may already at the outset possess some exogenous (imperfect) private information about their valuations.

In this context, when the information possessed by and provided to agents is private, the literature has shown that agents can secure information rents leading to distortions and welfare losses.[1] The main result of this paper shows that, to the contrary, in a large class of cases the principal can design information and a mechanism so as to fully extract the complete information first-best surplus. Hence, information design when coupled with mechanism design renders the privacy of both ex ante and ex post information entirely irrelevant!

The intuitive idea behind this result is that the principal can reduce information rents by concealing the information structure she uses to inform agents. This allows her to elicit agents'

[1] Most notably, this is implied by Li and Shi (2017a), Esö and Szentes (2007a,b), Krähmer and Strausz (2015a). The literature is reviewed in more detail below.

private information by cross-checking whether the information they report is consistent with the true information structure. Thus, my analysis highlights that information design may not only serve the purpose to inform but also to monitor agents.

More specifically, I distinguish two cases. I say that the principal has *full informational control* if the information that she can disclose pins down agents' valuations. In this case, an agent's initial private information only affects his beliefs about, yet does not directly influence, his valuation.[2] Otherwise, I say that the principal has only *partial informational control*.[3] I show that with full informational control, the principal can always extract the full complete information first-best surplus. With partial informational control, she can do so provided the correlation between an agent's initial information and the information the principal can disclose satisfies a "spanning" (or "full rank") condition as in the principal agent literature with contractible ex post information (Riordan and Sappington, 1988).[4]

While my results apply to general mechanism design settings with multiple agents, the underlying logic becomes clearest in the case with a single agent. Thus, the first part of the paper focusses on this case. For concreteness, I consider a seller (principal) who can design a sales mechanism and any information structure (such as a product sample) that provides the buyer (agent) with signals (for instance, taste experiences) which are informative about an otherwise unknown state that affects his valuation for the product (such as an unknown product feature that is of interest to the agent).

The novelty of my approach is the *combination* of two features: (i) I allow the principal to design various information structures and (to commit) to secretly randomize between them. This formalizes the above-mentioned idea that the principal uses the information structure as a moni-

[2]For example, in a sales context, when the seller introduces a new model variety, the buyer's experience with the old model may give him some (private) idea about, but not directly affect, his ultimate utility from the new model. Full informational control is assumed in Li and Shi (2017).

[3]For example, a buyer's valuation for a house may depend also on the (privately known) number of friends in the neighbourhood, and even if the principal fully discloses all features of the house, this does not pin down the buyer's valuation. Partial informational control is assumed in Esö and Szentes (2007a,b).

[4]Full surplus extraction mechanisms that exploit correlation are often deemed unrealistic because they violate real world constraints, such as ex post participation or cash constraints, or are considered too detail dependent. While I abstract from these constraints to keep the analysis clean, the basic force of employing an information structure as a monitoring device is likely to have *some* benefits also in environments in which the requirements for full surplus extraction are not exactly met.

toring device. One way to think about this is that the principal can "frame" the selling environment in which information disclosure takes place. For example, a car dealer may offer test drives, secretly employing various types of tires that affect the agent's driving experience (e.g. standard, sporty, comfortable, etc). An agent who cannot distinguish between the various tires is then uncertain whether his driving experience is due to the car as such or the tires. (ii) In the single agent case, I allow the parties to employ rich contracting protocols and to condition the terms of trade on (reports about) the outcome of the principal's randomization, that is, the actual information structure. (As will become clear below, allowing for rich contracting protocols, is not needed in the setting with multiple agents, however.)[5]

Allowing for these two features has two implications. First, the principal can elicit the private signal she supplies to the agent at no cost. The idea is to have the principal randomize over a set of information structures with the property that any signal that the agent may observe can be generated only by a subset of, yet not by all, possible information structures. In particular, if the agent reports a signal that cannot be generated by the realized information structure, it becomes apparent that he must have lied. In fact, I present a construction so that the agent believes that a deviation from truth-telling will be detected as a lie with positive probability. Truth-telling can then be induced by penalizing the agent if a lie is detected.

Second, the principal can allocate the product efficiently. In my construction, the principal will randomize over information structures which are each fully informative: knowing the signal the agent observes and knowing the true information structure reveals the true state. Hence, once the agent's signal is elicited and the information structure is verified, the state is revealed. As a consequence, under full informational control where the agent's valuation depends only on the state, the product can be allocated efficiently and the agent can be charged his valuation. With partial informational control, where the agent's valuation depends also on his initial private information, efficiency requires to also elicit this information. The key insight is that since the state is revealed ex post, the mechanism can de facto condition on the true state, as if the state was an ex post verifiable signal. Hence, if the above-mentioned spanning condition holds, it can

[5](i) and (ii) imply that the principal or a trusted mediator has private information about the realized information structure when the agent learns and reports his signal. While this possibility extends the conventional notion of information design, again, in the setting with multiple agents, I provide a construction where no party has private information about the information structure.

be extracted at no cost as in Riordan and Sappington (1988).

The logic from the single agent case essentially carries over to the general mechanism design setting with multiple agents (possibly with interdependent valuations). However, my full surplus extraction result with multiple agents is economically more significant because it holds even if the principal faces two additional constraints: First, an agent can only be informed about his own preferences, yet not about the preferences of other agents. This restriction respects the notion that, for example, in a private values auction, a bidder's valuation is his private information not only vis-a-vis the auctioneer, but also vis-a-vis the other bidders.[6] Second, I show that with multiple agents, it not necessary to employ rich contracting protocols that condition the terms of trade on the realized information structure: it is sufficient to use standard mechanisms which only condition on agents' reports about their private information. Moreover, while the principal still randomizes over information structure, it is not necessary that she (or any other player) observes its realization. Thus, the construction is entirely in line with standard notions of information and mechanism design.

The intuitive idea is that the principal uses various "salesmen" to inform agents about their valuations, for example, by offering product descriptions. While every salesman is truthful, they use a more or less "inflated" language to describe the product. At the outset, one salesman is picked at random that informs all agents about their valuations. Hence, when an agent receives a, say, very exaggerated product description, he infers that also other agents are likely to have received descriptions in an exaggerated range. Consequently, when reporting back an "understated" description, this is likely to be inconsistent with the other agents' reports, and hence detected as a lie. I show how this idea can be formalized to elicit signals truthfully and to identify the true state.

Finally, I investigate to what extent the spanning condition which is sufficient for full surplus extraction with partial informational control is also necessary. I show that if the spanning condition is violated, then there are always agent valuations so that for no information structure, full extraction of the complete information first-best surplus is feasible.[7] This includes information structures that depend on a report by an agent about his initial information (termed "discrimi-

[6]In addition, allowing the principal to inform agents about other agents' preferences makes the problem less interesting. For, the principal could then simply make all agents' preferences common knowledge among the agents and readily elicit them through some "shoot-the-liar" type of scheme.

[7]The argument is essentially identical to that of the "necessity" statement in Cremer and McLean (1988).

natory information disclosure" by Li and Shi, 2017a). However, if the information structure can depend not only on a report by an agent but also on his true initial information (termed "general disclosure" by Li and Shi, 2017b), then full surplus extraction of the complete information first-best surplus becomes again possible, even for all beliefs of the agents.

Related literature

The question I address in this paper is at the heart of a recent literature that studies information disclosure in mechanism design (or screening) where the principal, without observing herself, controls the additional private information agents learn beyond their initial private information.[8] My framework encompasses cases both with full informational control when the additional information provided by the principal and the agents' initial information is correlated, as in Li and Shi (2017a), or is orthogonal, as in Bergemann and Pesendorfer (2007), as well as with partial informational control when the additional and the initial information is orthogonal, as in Esö and Szentes (2007a,b).[9] While I show that the first-best is attainable in Li and Shi (2017a) and Bergemann and Pesendorfer (2007) type settings[10], I also show that this is generally not the case in Esö and Szentes (2007a,b) type settings, because if the initial and additional information are orthogonal, then the spanning condition mentioned earlier is violated.[11]

[8]For design settings where the additional information cannot be controlled by the principal, see, e.g., Baron and Besanko (1984) or Courty and Li (2000). See Krähmer and Strausz (2015b) for an overview.

[9]To be precise, Esö and Szentes (2007a,b) arrive at and work with such a model after applying their "orthogonalization" approach to a model in which the valuation is fully pinned down by the state. As Li and Shi (2017a) make clear, when the principal controls the additional information disclosed to an agent, the orthogonalization is not innocuous, because it matters whether the principal can disclose information about an agent's valuation or only about an orthogonal component of it. Moreover, my framework does not literally nest Li and Shi (2017a), Bergemann and Pesendorfer (2007), and Esö and Szentes (2007a,b) because to facilitate tractability, I only allow for discrete information.

[10]Bergemann and Wambach (2015) show that in the setting of Bergemann and Pesendorfer (2007), the first-best is attainable when the additional information is gradually disclosed and elicited.

[11]Li and Shi (2017a) do provide an example in which the principal is able to extract the full first-best surplus, but next to exploiting a particular distributional specification, the example rests on the unit good assumption, while my results hold also in the non-unit good case. Moreover, Li and Shi (2017a) show that the principal may benefit from using partial information disclosure as a price discrimination device. In my setting with richer contracting possibilities, the principal does not need to engage in discriminatory information disclosure, and the information disclosed to the agent is, by itself, only partially informative, but jointly with knowledge about the information structure reveals the agent's valuation.

My paper is closely related to recent work by Zhu (2018) who also considers an information plus mechanism design setting where agents have initial private information, and additional private information about an unknown state can be disclosed. When there are at least three agents whose preferences are linear with respect to the allocation (but possibly absent transfers), Zhu (2018) shows that the designer can implement the same outcome as in a benchmark in which agents only know their types, and the state is publicly revealed ex post. While my and Zhu's (2018) constructions are similar in spirit, the problem in Zhu (2018) originates from the constraint that agents have to report initial and additional private information simultaneously, rather than from the constraint that agents cannot be informed about others' preferences, as in my case. With two agents, Zhu (2018) shows that the benchmark can be attained when utility is quasi-linear and the state is orthogonal to the agents' initial private information. In contrast, while I only consider quasi-linear utility, my analysis covers the case with correlation, including the single agent case to which Zhu (2018) is not applicable. The most important difference is, however, that I show that not only the additional information provided to agents but also their initial private information can be elicited at no cost, and full surplus can be extracted.[12]

The idea that a designer can benefit from using random information structures is well-known from Myerson's (1982, 1986) work on mediation, and the more recent literature on information design (Bergemann and Morris, 2016) or Bayesian persuasion (Kamenica and Gentzkow, 2011). For, randomizing over information structures simply corresponds to the standard notion of a mediator randomizing over action recommendations for the agent(s) and is thus implicit in the appropriate notion of correlated equilibrium. What my paper makes clear is that in a framework where (some) actions are contractible and can condition on reports about the private signals provided to the agents, randomizing over information structures has the additional benefit that it facilitates the elicitation of these signals from the agents.

A similar point is also made in Rahman (2012) and Rahman and Obara (2010) who show that in team problems, making an agent's pay contingent on secret effort recommendations made to the other agents, fosters effort incentives and allows to elicit signals privately observed by an agent.[13] If the distribution of the private signal depends on others' efforts, making a secret (and incentive compatible) random effort recommendation to others corresponds to secretly random-

[12] I discuss the relation between Zhu's (2018) and my construction in more detail in Remark 4 and Appendix B.

[13] See also Strausz (2012).

izing over information structures in my setting, and making pay contingent on the effort recommendation made to others corresponds to making the terms of trade contingent on the outcome of the randomization in my setting. The difference is that in my set-up, the information structure can be freely designed, whereas in Rahman (2012) and Rahman and Obara (2010), the given team technology has to satisfy certain conditions for signals to be elicitable.

A different force is at work in Rodina (2018) who, in a career concerns framework with moral hazard, shows that making the information provided to the market contingent on a secret random effort recommendation to the agent may increase this agent's effort incentive. The agent will then hold different beliefs about the distribution of the market wage, depending on the recommendation, and this relaxes the effort constraint.[14]

Finally, the design of additional information for a privately informed agent is also considered in Bergemann, Bonnati, and Smolin (2017), but in their setting only transfers are contractible, and the agent takes a non-contractible action after information is revealed. While Bergemann, Bonnati, and Smolin (2017) restrict attention to simple contracts where the principal offers a menu of information structures and prices, allowing for the richer contracting protocols of my setting has the potential to improve the principal's revenue, but my first-best results do not directly carry over due to the presence of the additional obedience constraints resulting from the agent's non-contractible action.[15]

The paper is organized as follows. The next section presents the single agent case. Section 3 extends the analysis to multiple agents. Section 4 discusses extensions, and Section 5 concludes. All proofs are in the appendix.

[14]That a principal can benefit from endogenously creating correlation through randomization has also been observed in other contexts. Krähmer (2012) shows how an auctioneer can create correlation by randomizing over investments that improve bidders' valuations stochastically. In Obara (2008), bidders can take (hidden) actions that influence the joint distribution of their valuations, and almost full surplus extraction can be attained by a mechanism which implements a mixed action profile by bidders.

[15]Kolotilin et al. (2017) consider the design of information for a privately informed agent when both actions are non-contractible and monetary transfers are infeasible. In this case, there is no role for rich contracting protocols of the sort I study.

2 The model with a single agent

There is a principal (she) and an agent (he). The principal can produce a quantity $x \geq 0$ of some good at costs $c(x) \geq 0$ with $c(0) = 0$. The agent's valuation for consuming x may depend on two pieces of information, θ and ω.[16] It is common knowledge that θ is drawn from the set $\Theta = \{1, \ldots, \bar{\theta}\}$ with distribution $r \in \Delta(\Theta)$, and that, conditional on θ, ω is drawn from the set $\Omega = \{1, \ldots, \bar{\omega}\}$ with distribution $p_\theta \in \Delta(\Omega)$.[17] I impose the (mild) assumption that p_θ has full support Ω for all θ. I refer to θ as the agent's (ex ante) "type", and to ω as the "state". I denote the agent's valuation for quantity x by $v_{\theta\omega}(x) \geq 0$ with $v_{\theta\omega}(0) = 0$ for all θ, ω.[18]

The terms of trade consist of a quantity x and a payment t from the agent to the principal. The parties have quasi-linear utility, that is, if the terms of trade are x and t, the principal's utility is $t - c(x)$, and the agent's utility is $v_{\theta\omega}(x) - t$.

I assume that there is a well-defined first-best quantity given by

$$x^*_{\theta\omega} = \arg\max_x v_{\theta\omega}(x) - c(x), \tag{1}$$

and I denote the (expected) complete information first-best surplus by

$$Z^* = \sum_{\theta,\omega} r(\theta) p_\theta(\omega)[v_{\theta\omega}(x^*_{\theta\omega}) - c(x^*_{\theta\omega})]. \tag{2}$$

At the outset, the agent privately observes his type θ. In contrast, the state ω is not directly observable, neither by the agent nor the principal. However, the principal (and only the principal) can provide the agent with information about ω.[19] For example, the principal may offer product samples or give the agent more or less time to inspect and try out the product. More generally, I allow the principal to design any information structure that provides the agent with signals about the state. I assume that whatever the agent learns and infers from this information is not verifiable and the agent's private information.

[16] At the expense of more notation, all results go through essentially unchanged if also the principal's costs c depend on θ and ω.

[17] Assuming that Ω and Θ are finite is largely for simplicity of exposition. In Remark 6 below I show that if Ω is continuous, my construction can be extended to still achieve *approximate* full surplus extraction if $v_{\theta\omega}$ does not depend on θ, or if Θ is finite.

[18] The specification includes the frequently studied "unit good" case for: $v_{\theta\omega}(x) = v_{\theta\omega} x \cdot 1_{[0,1]}(x)$, and $c(x) = cx + \bar{c} \cdot 1_{(1,\infty)}(x)$ with $v_{\theta\omega} \geq 0$, $c \geq 0$ and $\bar{c} > 0$ large.

[19] Notice that the full support assumption rules out that (some type of) the agent knows ω for sure at the outset.

Formally, an information structure consists of a set S of signals and conditional signal distributions $\pi : \Omega \to \Delta(S)$. As it turns out the set of signals can be discrete to achieve full surplus extraction, and I set $S = \mathbb{Z}$ equal to the set of all integers.[20] I denote by $\pi_\omega(s)$ the conditional probability that signal s occurs, conditional on ω. Having fixed S, I refer to π as a ("simple") information structure.

Before I state the principal's problem formally, I introduce some important definitions. I say that the principal has *full informational control* if the agent's valuation only depends on the state: $v_{\theta\omega} = v_{\theta'\omega}$ for all θ, θ', ω. Note that this implies that the first-best quantity does not depend on the type: $x^*_{\theta\omega} = x^*_{\theta'\omega}$ for all θ, θ', ω. Otherwise, I say the principal has only *partial informational control*. I say that the agent's beliefs satisfy the *spanning condition* if there is no type $\tilde{\theta}$ whose belief $p_{\tilde{\theta}}$ can be written as a convex combination of the beliefs p_θ of the other types $\theta \neq \tilde{\theta}$.

Remark 1. With full informational control, the model corresponds to a discrete type version of Li and Shi (2017a,b). With partial informational control, and if types are orthogonal to states, that is, $p_\theta = p_{\tilde{\theta}}$ for all $\theta, \tilde{\theta}$, the model corresponds to a discrete type version of Esö and Szentes (2007a,b) (after their orthogonalization). Clearly, the spanning condition is violated in this case.

The principal's problem

The principal's objective is to design an information structure and a mechanism so as to maximize her expected payoff. The novelty of my approach is the *combination* of two features: I allow the principal to randomize among information structures and to employ "rich contracting protocols" which condition the terms of trade on the realized information structure. I focus on the case that the principal can randomize over at most countably many information structures. Formally, let $K = \mathbb{Z}$,[21] and let π_k be an information structure for $k \in K$, with $\pi_{\omega k}(s)$ denoting the conditional probability that signal s is observed, conditional on ω and k. The principal may (commit to) select information structures according to any distribution $\mu \in \Delta(K)$ where $\mu(k)$ is the probability with which π_k is selected. I denote the resulting ("compound") information structure by (Π, μ).

In addition, the principal designs a mechanism that specifies the terms of trade. Under a "rich

[20] Setting $S = \mathbb{Z}$ is primarily for ease of exposition. In Remark 4 below, I show that all my results can be established within a framework of finitely many signals. For the impossibility result in Proposition 3 below, I will allow for fully general information structures.

[21] Again, setting $K = \mathbb{Z}$ is primarily for ease of exposition. See Remark 4 below.

contracting protocol", the mechanism can condition on the realized information structure k.[22] This means, the information structure is verifiable ex post. Below I show that when there are multiple agents, my results go through for standard mechanisms that exclusively condition on communication by the agents, and information structures need not be verifiable. Similarly, I will argue that in the single agent case, my results go through if the mechanism can condition only on a report by the principal about k, and the parties can employ a budget breaker (see Remark 3).

The relationship between the principal and the agent proceeds as follows.

1. The agent privately observes θ.

2. The principal commits to an information structure (Π, μ) and a mechanism.

3. The agent decides to accept or reject.

   – If the agent rejects, every party gets their outside option of 0.

4. If the agent accepts, π_k is selected with probability $\mu(k)$, unobserved by the agent; and the agent privately observes a signal s generated by the information structure π_k.

5. The terms of trade are enforced according to the mechanism.

For a given information structure, the revelation principle (Myerson, 1986) implies that an optimal mechanism is in the class of direct and incentive compatible mechanisms which require the agent to submit a report $\hat{\theta}$ about his ex ante type after stage 3 and a report $\hat{s}$ about the signal observed after stage 4. I refer to $\hat{\theta}$ as an ex ante report and $\hat{s}$ as an ex post report.[23] Consequently, as I allow the terms of trade to condition on the realized information structure k, a mechanism consists of contingent quantities $x : \Theta \times S \times K \to \mathbb{R}_+$ and contingent transfers $t : \Theta \times S \times K \to \mathbb{R}$,

[22] To be sure, a compound information structure (Π, μ) induces a distribution over signals which corresponds to a "simple" information structure π where $\pi_\omega(s) = \sum_k \mu(k)\pi_{\omega k}(s)$. Therefore, the assumption that the principal can commit to a probability distribution over (simple) information structures corresponds to the standard assumption in the information design literature that the principal can use *any* (simple) information structure that provides signals to the agent. On the other hand, what distinguishes a simple and a compound information structure in my approach is that the principal or a trusted mediator privately learn the outcome of the realization so that contracts can condition on (reports about) it.

[23] It is common in the literature (notably Li and Shi 2017a), to also allow the information structure to depend on a report by the agent about his type, or even on the true type. I will discuss these cases in Section 4.

where $x(\hat{\theta},\hat{s},k)$ (resp. $t(\hat{\theta},\hat{s},k)$) denotes the quantity produced (resp. transfer paid) if the agent reports $\hat{\theta}$ and $\hat{s}$, and the realized information structure is π_k.

To express the principal's problem formally, I denote by $u(\theta,s;\hat{\theta},\hat{s})$ agent type θ's expected utility from reporting $\hat{s}$ ex post, conditional on having reported $\hat{\theta}$ ex ante and having observed s ex post (provided the probability of (θ,s) is positive). Moreover, let $U_{\theta,\hat{\theta}}$ be the expected utility of agent type θ from reporting $\hat{\theta}$ ex ante, that is,

$$U_{\theta,\hat{\theta}} = \sum_{\omega,k,s} p_\theta(\omega)\mu(k)\pi_{\omega k}(s)\max_{\hat{s}}[v_{\theta\omega}(x(\hat{\theta},\hat{s},k)) - t(\hat{\theta},\hat{s},k)]. \tag{3}$$

Finally, the principal's expected utility from information structure (Π,μ) and a mechanism (x,t) can be written as

$$W = \sum_{\theta,\omega,k,s} r(\theta)p_\theta(\omega)\mu(k)\pi_{\omega k}(s)[t(\theta,s,k) - c(x(\theta,s,k))], \tag{4}$$

and the principal's problem is given by

$$P: \quad \max_{(\Pi,\mu),(x,t)} W \quad s.t. \tag{5}$$

$$u(\theta,s;\theta,s) \geq u(\theta,s;\theta,\hat{s}) \quad \forall \theta,s,\hat{s}, \tag{6}$$

$$U_{\theta,\theta} \geq U_{\theta,\hat{\theta}} \quad \forall \theta,\hat{\theta}, \tag{7}$$

$$U_{\theta,\theta} \geq 0 \quad \forall \theta. \tag{8}$$

The first constraint is referred to as the ex post incentive compatibility constraint which ensures that the agent reports the signal truthfully ex post. Notice that the revelation principle requires truthful reporting of the signal only "on the path", that is, after a truthful ex ante report. The second constraint is the ex ante incentive compatibility constraint which ensures that the agent reports his type truthfully ex ante. The third constraint is the individual rationality constraint which ensures that all types accept the mechanism.[24] (I will argue in Remark 2 below that in the case with full informational control, a stronger individual rationality constraint holds which requires the agent to accept the mechanism only after having observed s.)

An information structure

I shall now define the information structure that I use below to construct full surplus extracting

[24]As usual, since the agent's outside option can be replicated in the mechanism by producing and charging nothing, it is without loss of generality optimal for the principal to induce all types to accept the mechanism.

first-best mechanisms. Let π_k^0 be defined by

$$\pi_{\omega k}^0(s) = \begin{cases} 1 & if \quad s = \omega + k, \\ 0 & else \qquad . \end{cases} \tag{9}$$

Under information structure Π^0, if the state is ω and the information structure is π_k^0, then the signal $s = \omega + k$ is released with probability 1. This has two important implications: First, knowing s and k reveals that the true state is $\omega = s - k$. Second, suppose that agent type θ has observed signal s. Then he assigns positive probability to k having occurred only if

$$s - k \in \Omega. \tag{10}$$

I say that s and k are *consistent* with one another if $s - k \in \Omega$, whereas they are *inconsistent* otherwise. Clearly, if the agent's ex post report $\hat{s}$ deviates from truth-telling, this is "detected" as a lie if the true k turns out to be inconsistent with the reported $\hat{s}$. The next lemma, while straightforward, is key to my results.

Lemma 1. *(i) There is μ^0 and $\beta > 0$ so that for all θ, s and all k that are consistent with s, we have $Pr(k \mid \theta, s) > \beta$.*
(ii) For all signals s and all reports $\hat{s} \neq s$, there is a $\kappa \in K$ so that s is consistent with κ, but $\hat{s}$ is inconsistent with κ.

Part (i) says that if agent type θ has observed signal s, then he assigns to any k that is consistent with s a probability that is bounded from below when μ is chosen appropriately. Intuitively, this is so, because there is only a finite set of k's that are consistent with a given s. Part (ii) together with (i) implies that the agent expects any deviation from truth-telling to be detected as a lie with at least probability $\beta > 0$. As a consequence, by penalizing reports sufficiently harshly if they are detected as a lie, the agent is induced to report his signal truthfully. But once the signal is elicited truthfully, the signal together with the information structure identifies the true state. As I show next, this can be used to extract the full surplus, which is the main result for the single agent case.

Full surplus extraction

Proposition 1. *Let Π^0 and μ^0 from Lemma 1 be given.*

(i) Suppose the principal has full informational control. Then there is a mechanism which implements the first-best, and the principal fully extracts the surplus Z^*.

(ii) *Suppose the principal has only partial informational control. If the agent's beliefs satisfy the spanning condition, then there is a mechanism which implements the first-best, and the principal fully extracts the surplus* Z^*.

The intuition behind the proposition is as follows. By Lemma 1, the signal s can be elicited at no cost from the agent by penalizing reports that are inconsistent with the realized information structure. But if the signal is truthfully elicited, then, since s and k reveal that the true state is $\omega = s - k$, this means that the mechanism can effectively condition on the true state directly.

If the principal has full informational control, it is therefore as if the agent has no (payoff relevant) private information at all, and one can simply implement the first-best quantity x^*_ω and charge the agent his valuation $v_\omega(x^*_\omega)$ if state ω is revealed. This mechanism attains the first-best and extracts the full surplus.

In contrast, if the principal has only partial informational control, the mechanism needs to elicit the agent's type θ to attain the first-best. The crucial observation is the following. Because the state ω is fully revealed once s and k are known, the state is de facto verifiable, and therefore it may also serve the role of an ex post verifiable signal. But if the principal has access to such a signal, and the signal is correlated with the agent's type in the sense that the spanning condition holds, then she can elicit the agent's type at no cost which is well-known from the literature on principal agent problems with ex post verifiable information (Riordan and Sappington, 1988).

Before I turn to the multiple agent case, I close this section with 3 remarks that show extensions and corollaries of Proposition 1.

Remark 2 (Individual rationality)**.** As the paragraph preceding the previous paragraph makes clear (and as the proof of Proposition 1 shows), in the case with full informational control, the mechanism satisfies a stronger individual rationality constraint than (8). In fact, for all s (resp. all ω), the mechanism delivers the agent zero utility conditional on s (resp. on ω). This means that the mechanism is individually rational even if the agent can still choose his outside option after having observed the signal s. Moreover, "on path", the agent does not make a loss ex post after ω has been revealed. Finally, the mechanism does not elicit θ. This means that, unlike in Li/Shi (2017a), the mechanism does not screen the agent sequentially.

Remark 3 (Non-verifiable information structure)**.** Proposition 1 rests on the fact that the information structure k is verifiable so that the mechanism can directly condition on it. This assumption can be dropped when the principal privately observes k and is required to make a report about the information structure k. That is, after stage 4 in the time-line above, the principal and the agent simultaneously report $\hat{k}$ and $\hat{s}$ respectively. The basic idea is to induce truth-telling by cross-checking the parties' reports and penalizing both parties if their reports $\hat{s}$ and $\hat{k}$ are inconsistent with one another, that is, $\hat{s} - \hat{k} \notin \Omega$. Given the principal reports truthfully, truth-telling is a best response for the agent for the same reasons as in the case with verifiable k. A similar logic applies to the principal. Because the mechanism penalizes both the agent and the principal if inconsistent reports are submitted, it is not budget-balanced off the path and requires a third party to cash the penalty. Because the logic of the construction is essentially the same as the logic behind Proposition 2 for the case with multiple agents, I omit further details.

Remark 4 (Finite information structures)**.** The information structure (Π^0, μ^0) uses countably many signals and countably many information structures π_k^0. Similar to ideas in Zhu (2018), I now briefly illustrate how full surplus can be extracted using only finite S and finite K.[25] Recall that the number of states is $\bar{\omega}$. Let $S = K = \{0, 1, \ldots, \bar{\omega}\}$, and conditional on ω and k, let the signal that is released with probability 1 be equal to $s = \omega + k$ modulo $\bar{\omega} + 1$.[26] To illustrate, consider the case with three states. The following table depicts the signal s that is released (with probability 1), conditional on ω and k:

s	$k=0$	$k=1$	$k=2$	$k=3$
$\omega=1$	1	2	3	0
$\omega=2$	2	3	0	1
$\omega=3$	3	0	1	2

Table 1: Information structure with finitely many signals

It can be seen by inspection that s and k identify the true ω. Moreover, Lemma 1 holds: Signal s is consistent with all $k \neq s$. A lie $\hat{s} \neq s$ is inconsistent with $k = \hat{s}$, and because Θ, S, K are finite, we have for all μ with full support that $Pr(k \mid \theta, s)$ is bounded from below for all k that are consistent with s.[27]

[25] I discuss the relation between my work and Zhu (2018) in more detail in Appendix B. I also thank Nicolas Schutz for a discussion of these ideas.

[26] Recall that for two natural numbers m, n, m modulo n is the remainder of m/n.

[27] In contrast to my original construction (9), the "meaning" of signals in this information structure is not "mono-

Remark 5 (Unit good)**.** A setting that is prominently discussed in the literature is the "unit good case". In particular, Li and Shi (2017a) consider the unit good case with full informational control. I now show that in this case full surplus can be extracted with only three simple information structures and three signals. Indeed, in the unit good case with full informational control, we have $x \in [0,1]$, the buyer's valuation is given by $v_\omega x \geq 0$, and the seller's cost is $cx \geq 0$. Let $\Omega^+ = \{\omega \mid v_\omega \geq c\}$ (resp. $\Omega^- = \{\omega \mid v_\omega < c\}$) be the set of states in which consumption is (resp. is not) efficient. Because all that matters for efficiency is whether ω is in Ω^+ or in Ω^-, this effectively corresponds to the case with two states. In the spirit of Remark 4, consider the information structure:

s	$k=0$	$k=1$	$k=2$
$\omega \in \Omega^+$	1	2	0
$\omega \in \Omega^-$	2	0	1

Table 2: Information structure with 3 signals for the unit good case

The information structure identifies whether consumption is efficient or not (e.g., $s=1$ and $k=0$ identify that it is, while $s=1$ and $k=2$ identify that it is not). As above, truth-telling of s can be induced by penalizing inconsistent reports. The following mechanism then extracts full surplus: if the agent's report s and k identify a state in Ω^+, then the agent gets the good and is charged his conditional expected valuation $E[v_\omega \mid \omega \in \Omega^+]$. Otherwise, he does not get the good and is charged nothing. Note also that the mechanism is individually rational, conditional on s, and does not elicit θ. Moreover, the argument does not depend on the discreteness of Ω or Θ and goes trough unchanged if both spaces are continuous (as long as p_θ has support Ω).

Remark 6 (Continuous state and type space)**.** If the state space in continuous, say $\Omega = [\underline{\omega}, \bar{\omega}]$, $\underline{\omega} < \bar{\omega}$, then my construction can be adapted by discretizing Ω. More formally, consider a partition of Ω in segments $\Omega_\ell = [\omega_\ell, \omega_{\ell+1})$, $\ell = 1, \ldots, N$, $\omega_1 = \underline{\omega}$, $\omega_{N+1} = \bar{\omega}$, $\omega_\ell < \omega_{\ell+1}$. Consider now the (compound) information structure where the principal randomizes over $k \in \mathbb{Z}$, and the signal $s = \ell + k$ is revealed with probability 1 conditional on k and $\omega \in \Omega_\ell$. The construction used for Proposition 1 can then be used to elicit without cost whether $\omega \in \Omega_\ell$. Under full informational

tone". That is, numerically larger signals do not always indicated numerically larger states (e.g., when $k=1$, $s=3$ reveals that the state is $\omega = 2$ while $s=0$ reveals $\omega = 3$). When signals correspond to taste experiences from product samples, this is not straightforward to interpret.

control, this implies that by choosing the partition sufficiently finely, the principal can attain the first best surplus almost fully (irrespective of whether the state space Θ is discrete or continuous). This construction also permits almost full surplus extraction under partial informational control if Θ is finite (note that the spanning condition now depends on how Ω is partitioned and can be made to hold by design of the partition). If Θ is continuous, then almost full surplus extraction can be attained under the belief conditions in McAfee and Reny (1992).

3 Multiple agents

In this section, I show that Proposition 1 extends to a general mechanism design setting with multiple agents. While the underlying logic is analogous to the single agent case, I stress two important economic differences: first, in the multiple agents case, full surplus can be extracted with standard mechanisms that condition exclusively on communication by the agents but do not condition on the realized information structure itself. In particular, it is not needed that the realized information structure be verifiable and/or the principal is privately informed about it. Second, while in the single agent case, the principal (or the mediator) privately observes the realized information structure, I present a construction in which no player privately knows the realized information structure. Thus, the construction is entirely in line with standard notions of information and mechanism design.

In addition, with multiple agents the question arises if the principal can inform an agent about the preferences of the other agents. While this might be not implausible in some contexts, for reasons explained shortly, I will impose the constraint that the principal can provide an agent with information about this agent's yet not another agent's preferences.

Formally, there are now (for notational simplicity only) two agents, $i = 1, 2$. Let $\theta_i \in \Theta_i = \{1, \ldots, \bar{\theta}_i\}$ be agent i's type, and let $\Theta = \Theta_1 \times \Theta_2$ be the set of type profiles with generic element $\theta = (\theta_1, \theta_2)$. Let $\omega_i \in \Omega_i = \{1, \ldots, \bar{\omega}_i\}$ be "agent i's state", and let $\Omega = \Omega_1 \times \Omega_2$ be the set of "states" with generic element $\omega = (\omega_1, \omega_2)$. I assume that (θ_1, ω_1) and (θ_2, ω_2) are stochastically independent.[28] Let $r_i \in \Delta(\Theta_i)$ be the prior distribution of agent i's type, and let $p^{(i)}_{\theta_i} \in \Delta(\Omega_i)$ be type θ_i's belief about ω_i. I assume that $p^{(i)}_{\theta_i}$ has full support on Ω_i for all $\theta_i \in \Theta_i$, $i = 1, 2$.

[28] If (θ_1, ω_1) and (θ_2, ω_2) were correlated across agents, the problem would be less interesting, because full surplus extraction would (often) immediately follow from Cremer and McLean (1988).

The terms of trade consist of a contractible allocation x in some set X of allocations and of transfers $t_i \in \mathbb{R}$ from the agents to the principal.[29] The parties have quasi-linear utility: given θ and ω, the principal's utility from the terms of trade (x, t_1, t_2) is denoted $v^P(x) + t_1 + t_2$, and agent i's utility is $v^{(i)}_{\theta,\omega}(x) - t_i$. Thus, I allow for interdependent valuations among the agents.[30] To abstract from complications that may arise from individual rationality, I assume that there is the option to exclude an agent from the mechanism without affecting the other players. Formally, there is an allocation x^0 which yields any party a gross utility of 0, and for any allocation x and each i, there is an allocation x^0_i so that agent i gets 0 from x^0_i, and the principal and agent $j \neq i$ get the same as under allocation x.

I assume that there is a well-defined first-best allocation given by

$$x^*_{\theta,\omega} = \arg\max_{x \in X} \left[\sum_{i=1,2} v^{(i)}_{\theta,\omega}(x) + v^P(x) \right], \tag{11}$$

and I denote the (expected) complete information first-best surplus by

$$Z^* = E\left[\sum_{i=1,2} v^{(i)}_{\theta,\omega}(x^*_{\theta,\omega}) + v^P(x^*_{\theta,\omega}) \right]. \tag{12}$$

At the outset, each agent i privately observes his type θ_i, but no agent can observe the state ω, and only the principal can provide information about ω. Whatever an agent learns and infers from this information is not verifiable and this agent's private information. Importantly, I impose the above-mentioned information disclosure constraint on the principal:

Constraint C *The principal can provide agent i with information about* ω_i *only, yet not about* ω_j, $j \neq i$.

Constraint C captures that an agent's preferences are ultimately this agent's private information not only vis-a-vis the principal but also vis-a-vis the other agent. For example, in a private values auction, where agents' valuations reflect their idiosyncratic tastes, by offering test samples, or by varying the time an agent is allowed to inspect the object, a principal can influence how well an agent is informed about his valuation, yet in doing so, cannot affect what an agent believes about the other agent's valuation. A second justification is that, absent constraint C, the

[29]This encompasses private goods settings. For example, in a unit good auction $X = [0,1]^2$, with $x_1 + x_2 \leq 1$, where x_i is the probability that agent i gets the object.

[30]As in the single agent case, at the expense of more notation, all results readily extend to the case that the principal's utility depends on θ and ω.

principal's problem becomes almost trivial. She could simply make ω common knowledge among the agents and then elicit it through some type of "shoot-the-liar" scheme.

I say that the principal has *full informational control* if for all i, x, we have:

$$E[v^{(i)}_{\theta,\omega}(x) \mid \theta_i, \omega_i] = E[v^{(i)}_{\theta,\omega}(x) \mid \omega_i] \qquad \forall \theta_i, \omega_i, \tag{13}$$

$$x^*_{\theta,\omega} = x^*_{\omega} \qquad \forall \theta, \omega. \tag{14}$$

Property (13) means that agent i's (expected) valuation for allocation x depends only on ω_i, not on his type θ_i. Property (14) says that the first-best allocation is independent of types.[31] If one of the properties is violated, I say the principal has only *partial informational control*. I say that agent i's beliefs satisfy the *spanning condition* if there is no type $\tilde{\theta}_i$ whose belief $p^{(i)}_{\tilde{\theta}_i}$ can be written as a convex combination of the beliefs $p^{(i)}_{\theta_i}$ of the other types $\theta_i \neq \tilde{\theta}_i$.[32]

I shall now discuss two approaches to extract full surplus. The first approach is a straightforward extension of the single agent case and shares with it that some players privately observe the realization of the information structure. In contrast, in the second approach illustrated in Remark 7, the realized information structure is not observed by anyone.

As for the first approach, let $K_i = \mathbb{Z}$, $i = 1, 2$. Let $K = K_1 \times K_2$, and take a probability distribution $\mu \in \Delta(K)$.Let

$$\sigma_i = \omega_i + k_i. \tag{15}$$

Given k and ω, agent i privately observes the signal

$$s_i = (\sigma_i, k_j), \quad j \neq i \tag{16}$$

with probability 1. In other words, agent i observes a noisy signal σ_i of his state ω_i and, in addition, is informed about the ("marginal") information structure k_j that is used for the other agent $j \neq i$. I denote the space of signals for agent i by $S_i = \mathbb{Z}^2$, and let $S = S_1 \times S_2$. Clearly, the information structure (16) satisfies constraint C.

[31] Note that in general neither of the properties implies the other.

[32] When the principal has full informational control, and types are orthogonal to states, then agents have no relevant ex ante private information, and my model essentially corresponds to Bergemann and Pesendorfer (2007). The former, however, explicitly imposes the participation constraint that agents can reject the mechanism after signals have been disclosed. Analogous arguments as in Remark 2 can be used to see that this constraint will be also be met in my setting by the optimal mechanism established below.

A mechanism now induces an extensive game between the agents with the analogous timing as in the single agent case, and I assume that agents play a Perfect Bayesian Equilibrium (PBE). Appealing again to the revelation principle, given the information structure (16), I can restrict attention to direct and (Bayesian) incentive compatible mechanisms which require each agent i to simultaneously submit a report $\hat{\theta}_i \in \Theta_i$ about his type after the contracting stage (after stage 3) and a report $\hat{s}_i \in S_i$ about the signal observed (after stage 4). I refer to $\hat{\theta}_i$ as an ex ante report and to $\hat{s}_i$ as an ex post report. A mechanism is incentive compatible if, given truth-telling by the other agent, each agent submits a truthful ex ante report and a truthful ex post report (conditional on having reported truthfully ex ante).

As indicated above, unlike in the single agent case, I now allow a mechanism to condition only on agents' reports about their private information, yet not on the true information structure k. I refer to such a mechanism as a "standard" mechanism, which, accordingly, consists of a contingent allocation $x : \Theta \times S \to X$ and contingent transfers $t_i : \Theta \times S \to \mathbb{R}$, $i = 1, 2$, from agent i to the principal.

I now show that there is a standard mechanism that guarantees full surplus extraction.

Proposition 2. *Suppose there are multiple agents, and let information structure (16) be given. There is a probability measure $\mu^0 \in \Delta(K)$ so that we have:*

(i) If the principal has full informational control, then there is a standard mechanism which implements the first-best, and the principal fully extracts the surplus Z^*.

(ii) If the principal has only partial informational control, and if all agents' beliefs satisfy the spanning condition, then there is a standard mechanism which implements the first-best, and the principal fully extracts the surplus Z^*.

The key step behind the proposition is to show that the signals $s_1 = (\sigma_1, k_2)$ and $s_2 = (\sigma_2, k_1)$ can be truthfully elicited from the agents at no cost. Once this is achieved, the state ω is known, because $\omega_i = \sigma_i - k_i$. As in the single agent case, if the principal has full informational control, the first-best allocation x^*_ω, which is independent of types by (14), can then be implemented, and any agent can be charged his (expected) valuation in the first-best, conditional on ω_i: $E[v^{(i)}_{\theta,\omega}(x^*_\omega) \mid \theta_i, \omega_i]$, which is independent of θ_i by (13). If the principal has only partial informational control and the spanning condition holds, then, as in the single agent case, the state ω_i associated with agent i can be used as a verifiable signal to elicit at no cost agent i's type θ_i.

The basic idea to elicit signals s_i is to cross-check agents' reports, and to penalize both agents if their reports $\hat{s}_1 = (\hat{\sigma}_1, \hat{k}_2)$ and $\hat{s}_2 = (\hat{\sigma}_2, \hat{k}_1)$ are inconsistent with one another, that is, if

$$\hat{\sigma}_\ell - \hat{k}_\ell \notin \Omega_\ell \quad \text{for some } \ell = 1, 2. \tag{17}$$

The logic of Lemma 1 can now be extended to show that truth-telling is an equilibrium. Suppose that agent i has observed s_i, and agent j reports s_j truthfully, $i \neq j$. Then (i) a probability measure μ^0 can be chosen so that agent i assigns positive probability, again bounded from below, to s_j having occurred if and only if s_i and s_j are consistent. And (ii), for any lie $\hat{s}_i \neq s_i$, there is a signal s_j so that s_i and s_j are consistent but $\hat{s}_i$ and s_j are inconsistent. Thus, agent i expects any deviation from truth-telling to be inconsistent with agent j's report with a probability that is bounded from below. Consequently, lying can be deterred by sufficiently penalizing inconsistent reports.

Remark 7. While there is nothing in the general notion of an information structure that prevents the principal from privately informing an agent about the (marginal) information structure used for the other agent, this may be difficult to implement in some applications. However, as I shall now briefly indicate, ideas analogous to Proposition 2 can be used to extract full surplus without any party obtaining private information about the realized information structure.

To illustrate the basic idea, consider the case with two states per agent: $\Omega_i = \{1, 2\}$, $i = 1, 2$, and let $K = \mathbb{Z}$. Suppose that conditional on ω and k, agents 1 and 2 respectively observe

$$s_1 = \omega_1 + k, \quad s_2 = \omega_2 + 1/2 \cdot k \tag{18}$$

with probability 1. Observe that constraint C is met.

One interpretation of (18) is that the principal uses "salesmen" to inform agents about their valuations, for example, by offering product descriptions. Each k corresponds to a different salesman, and while every salesman is truthful, the larger is k, the more "inflated" the language the salesman uses to describe the product.

The support of the joint distribution of (s_1, s_2) is illustrated in Figure 1. Each set of pairs (s_1, s_2) located on the corners of a square with a distinct dot shape corresponds to the support of the distribution of (s_1, s_2), conditional on some k. Because these supports are disjoint for all k, any pair (s_1, s_2) of signals identifies the true k and therefore both ω_1 and ω_2. (For the information structure to have this property is why k in the signal for agent 2 is scaled with $1/2$.) Moreover, as with Proposition 2, signals can be elicited from the agents without cost by punishing inconsistent

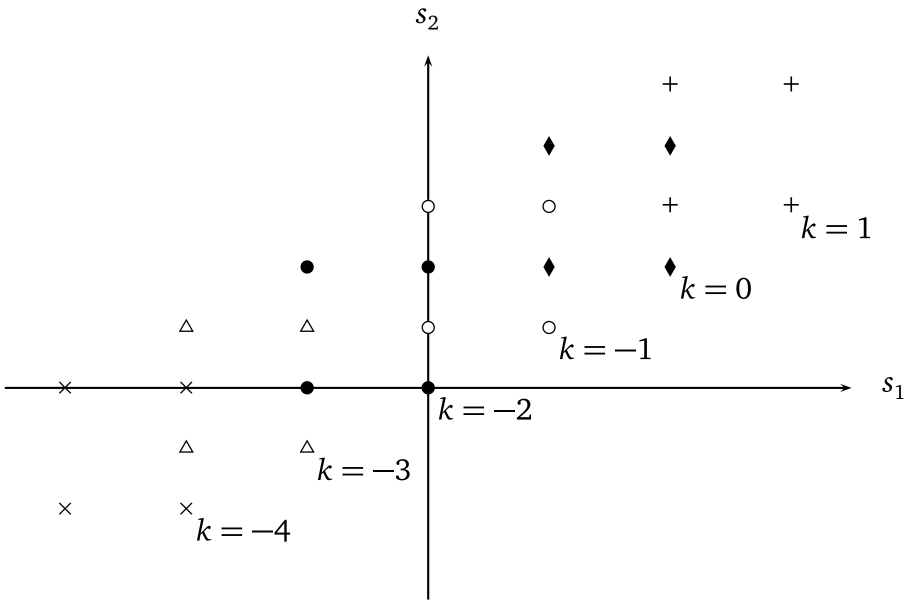

Figure 1: Support of (s_1, s_2) induced by (18)

reports, where signals s_1 and s_2 are now inconsistent if (s_1, s_2) is not in the support of the joint signal distribution. As above, it can be shown that given truth-telling by the other agent, each agent attaches a positive probability to a lie being inconsistent with the other agent's report. Hence, lies can be deterred by penalizing inconsistent reports sufficiently.

4 Contingent information structures

The information structures considered so far have the feature that they do not depend on ex ante reports by the agents about their types, or on these types themselves. In this section, I relax this feature to address the question to what extent the spanning condition is also necessary for full surplus extraction. To study this issue, I return to the single agent case. All arguments equally apply to the multiple agents case.

Recall the single agent case setting from section 2. I say, an information structure π_k is report-contingent if it depends on an ex ante report $\hat{\theta}$ by the agent, and I denote by $\pi_{\omega k}(s; \hat{\theta})$ the probability that a signal s is generated conditional on ω and k when the agent reports $\hat{\theta}$.[33] I say an information structure π_k is report- and type-contingent if it depends both on an ex ante report $\hat{\theta}$ by the agent and also on his true type θ.[34] I denote by $\pi_{\omega k}(s; \hat{\theta}, \theta)$ the probability that a signal

[33] Li and Shi (2017a) refer to this case as "discriminatory disclosure".

[34] See Li and Shi (2017b) for a discussion of this case when the principal cannot randomize among information

s is generated conditional on ω and k when the agent is of type θ and reports $\hat{\theta}$.[35]

I first show that allowing for report-contingent information structures does, in general, not help to relax the spanning condition. More precisely, I show that for a given set of beliefs that violates the spanning condition, there is a specification of the agent's valuation for which full surplus extraction is not possible. In this sense, the spanning condition is necessary for full surplus extraction. To make this impossibility result stronger, I now allow the sets S and K to be arbitrary measure spaces endowed with σ-algebras.

Proposition 3. *Suppose the principal has only partial informational control. If the agent's prior beliefs violate the spanning condition, then there are valuations $v_{\theta\omega}$ so that for any report-contingent information structure, any mechanism that implements the first-best leaves an information rent to some agent type θ.*

The argument is a straightforward adaptation of the analogous argument in the proof of Theorem 2 in Cremer and McLean (1988). The spanning condition is violated in the important class of settings considered by Esö and Szentes (2007a,b) where the type is orthogonal to the state so that p_θ does not depend on θ. Full surplus extraction is then not guaranteed.

Next, I allow for report- and type-contingent information structures. As shown next, this allows the principal to fully extract the first-best surplus irrespective of the agent's beliefs.

Proposition 4. *There is a report- and type-contingent information structure and a mechanism which implements the first-best, and the principal fully extracts the surplus Z^*.*

The intuition is straightforward. When the information structure can be conditioned on both a report about and the true type, then the principal can release, for all states ω and information structures k, the same signal $s_0 \in S$ ex post if the agent reports a type which differs from the true

structures and/or use rich contracting protocols.

[35]As an economic example for a report- and type-contingent information structure, consider a good that consists of various attributes $\theta \in \{1, \dots, \bar{\theta}\}$. The agent cares only about exactly one attribute θ which corresponds to his privately known type. The agent's valuation for the good depends in addition on an unknown state $\omega = (\omega_1, \dots, \omega_{\bar{\theta}})$, and is given by $v_{\theta\omega} = \phi(\omega_\theta, \theta)$ for some function ϕ. Now consider the following disclosure policy by the principal. If the agent announces $\hat{\theta}$, then some information about $\omega_{\hat{\theta}}$ is disclosed to the agent, but none about ω_θ, $\theta \neq \hat{\theta}$. Hence, if the agent announces $\hat{\theta}$, he receives information about his valuation if his true type is $\hat{\theta}$ but no information if his true type is $\theta \neq \hat{\theta}$. Hence, the information he receives depends on his true type. For an analysis of infomation disclosure and pricing with a multi-attributes good, see Smolin (2018).

type. This not only gives the agent no ex post information about the state, it also gives him no ex post information about the information structure k. The idea is again to impose a penalty on the agent if he reports a signal which is inconsistent with k. Hence, if the agent lies about θ, he will receive no additional information about ω and k, and so for any report $\hat{s}$, he expects to be penalized with positive probability. For sufficiently large penalty, this will deter the agent from misreporting ex ante. Thus, the agent's type can be elicited without leaving rents to the agent.

5 Conclusion

In this paper I show that in an information plus mechanism design setting with a single or multiple agents who possess imperfect private initial information, the principal can in a large class of cases design additional private information in such a way to fully extract the complete information first-best surplus. The basic idea is that information design can not only be used to inform but also to monitor agents: by secretly randomizing over information structures, the principal is put in a position where she can cross-check whether agents' reports are consistent with the realized information structure or other agents' reports. In this way, the private information provided to agents can be elicited without cost by penalizing inconsistent reports.

While the full surplus results shown in this paper are an extreme manifestation of the benefits of randomizing over information structures, it is likely that the basic drivers behind my results have force also in environments where full surplus extraction is not possible due to the presence of real world constraints such as, for example, ex post participation or cash constraints. It is an interesting avenue for future research to explore the benefits of randomizing over information structures when these constraints are explicitly modelled.

On a related note, by randomizing over information structures, the principal endogenously creates correlation among agents' valuations. An interesting question is whether this may enhance the principal's revenue also in settings where she cannot fine-tune the mechanism but is restricted to a specific format, such as first or second price auctions.

A Appendix

Proof of Lemma 1 (i) Let $\mu^0(k) = \gamma^{|k|+1}$ with γ such that $\sum_k \mu^0(k) = 1$.[36] Let

$$M = \min_{\theta}\left(\frac{\min_{\omega} p_{\theta}(\omega)}{\max_{\omega} p_{\theta}(\omega)}\right). \tag{19}$$

Note that $M > 0$ because $p_{\theta}(\omega) > 0$ for all θ, ω by assumption. Consider now first the case that $s - \bar{\omega} > 0$, and let k be consistent with s, that is, $k \in \{s - \bar{\omega}, \dots, s-1\}$. By Bayes' rule,

$$Pr(k \mid \theta, s) = \frac{\sum_{\omega} p_{\theta}(\omega)\pi^0_{\omega k}(s)\mu^0(k)}{\sum_{\omega,\ell} p_{\theta}(\omega)\pi^0_{\omega\ell}(s)\mu^0(\ell)} = \frac{p_{\theta}(s-k)\mu^0(k)}{\sum_{\ell=s-\bar{\omega}}^{s-1} p_{\theta}(s-\ell)\mu^0(\ell)} \geq M\frac{\gamma^{k+1}}{\sum_{\ell=s-\bar{\omega}}^{s-1}\gamma^{\ell+1}}, \tag{20}$$

where the equality uses the definition of Π^0, and the inequality uses the definition of μ^0 and M. Now observe that since $k \in \{s - \bar{\omega}, \dots, s-1\}$ and $s - \bar{\omega} > 0$,

$$\frac{\gamma^{k+1}}{\sum_{\ell=s-\bar{\omega}}^{s-1}\gamma^{\ell+1}} \geq \frac{\gamma^{s}}{\gamma^{s}\sum_{\ell=-\bar{\omega}}^{-1}\gamma^{\ell+1}} = \frac{1}{\sum_{\ell=-\bar{\omega}}^{-1}\gamma^{\ell+1}}. \tag{21}$$

Because this expression and M are independent of θ, s, k, it follows that $Pr(k \mid \theta, s)$ is bounded from below, as desired. The argument for the case $s - \bar{\omega} \leq 0$ is analogous.
(ii) Let $s, \hat{s}$ with $\hat{s} \neq s$ be given. If $\hat{s} < s$, then for $\kappa = s-1$, we have that $s - \kappa = 1 \in \Omega$ but $\hat{s} - \kappa < 1 \notin \Omega$. If $\hat{s} > s$, then for $\kappa = s - \bar{\omega}$, we have that $s - \kappa = \bar{\omega} \in \Omega$ but $\hat{s} - \kappa > \bar{\omega} \notin \Omega$. qed

Proof of Proposition 1 To define the mechanism, consider the following auxiliary problem over the choice variable $\tau : \Theta \times \Omega \to \mathbb{R}$:

$$\tilde{P}: \quad \max_{\tau(\theta,\omega)} \quad \sum_{\theta,\omega} r(\theta)p_{\theta}(\omega)[\tau(\theta,\omega) - c(x^*_{\theta\omega})] \quad s.t. \tag{22}$$

$$\sum_{\omega} p_{\theta}(\omega)[v_{\theta\omega}(x^*_{\theta\omega}) - \tau(\theta,\omega)] \geq \sum_{\omega} p_{\theta}(\omega)[v_{\theta\omega}(x^*_{\hat{\theta}\omega}) - \tau(\hat{\theta},\omega)] \quad \forall \theta, \hat{\theta}. \tag{23}$$

$$\sum_{\omega} p_{\theta}(\omega)[v_{\theta\omega}(x^*_{\theta\omega}) - \tau(\theta,\omega)] \geq 0 \quad \forall \theta. \tag{24}$$

Problem $\tilde{P}$ corresponds to the (static) principal agent problem, where θ is the agent's private information at the contracting stage, and the state ω is revealed ex post and is verifiable. Constraints (23) and (24) are the incentive compatibility and the individual rationality constraints. The principal's objective is to choose payments τ which can condition both on a report $\hat{\theta}$ and on the (ex post verifiable) state ω so as to maximize her revenue given the allocation rule $x^*_{\hat{\theta}\omega}$.

[36] γ is the solution in $(0,1)$ to the equation $\gamma^2 + 2\gamma - 1 = 0$.

Now, if the principal has full informational control so that $v_{\theta\omega}$ and $x^*_{\theta\omega}$ do not depend on θ, then a solution to problem $\tilde{P}$ is readily given by $\tau(\theta,\omega) = \tau(\omega) = v_\omega(x^*_\omega)$, because in this case, (23) is trivially satisfied, and (24) is binding. If the principal has only partial informational control, then, as shown by Riordan and Sappington (1988), there is a solution $\tau(\theta,\omega)$ to $\tilde{P}$ with (24) being binding if the beliefs p_θ satisfy the spanning condition. In either case, the value of the problem is the complete information first-best surplus Z^*.

I now define the mechanism (x,t) by

$$x(\theta,s,k) = \begin{cases} x^*_{\theta,s-k} & if \quad s-k \in \Omega \\ 0 & if \quad s-k \notin \Omega \end{cases}, \qquad t(\theta,s,k) = \begin{cases} \tau(\theta,s-k) & if \quad s-k \in \Omega \\ T & if \quad s-k \notin \Omega \end{cases} \tag{25}$$

for some $T > 0$.

I shall show that (x,t) is a solution to the original problem P that delivers Z^* to the principal. To prove this, I show that (x,t) satisfies (i) ex post incentive compatibility (6), (ii) ex ante incentive compatibility (7), (iii) individual rationality (8), and (iv) yields the principal Z^*.

To see (i), suppose agent type θ has reported $\hat{\theta}$ and observed s. The agent's utility from reporting $\hat{s} = s$ is

$$u(\theta,s;\hat{\theta},s) = \sum_{k:s-k\in\Omega} Pr(k \mid \theta,s)\{E_{\tilde{\omega}}[v_{\theta\tilde{\omega}}(x(\hat{\theta},s,k)) \mid \theta,s,k] - t(\hat{\theta},s,k)\}. \tag{26}$$

Note that this expression is independent of T, since the sum is only over indices k with $s-k \in \Omega$.

On the other hand, by Lemma 1, (ii), for any $\hat{s} \neq s$, there is κ so that $s-\kappa \in \Omega$ but $\hat{s}-\kappa \notin \Omega$. Thus, if the agent reports $\hat{s} \neq s$, then by part (i) of Lemma 1, with (at least) probability $Pr(\kappa \mid \theta,s) > b$, he has to make the payment

$$\tau(\hat{\theta},\hat{s}-\kappa) = T, \tag{27}$$

and his utility is

$$\begin{aligned} u(\theta,s;\hat{\theta},\hat{s}) \quad = \quad & \sum_{k:s-k\in\Omega,k\neq\kappa} Pr(k \mid \theta,s)\{E_{\tilde{\omega}}[v_{\theta\tilde{\omega}}(x(\hat{\theta},\hat{s},k)) \mid \theta,s,k] - t(\hat{\theta},\hat{s},k)\} \\ & + Pr(\kappa \mid \theta,s)\{E_{\tilde{\omega}}[v_{\theta\tilde{\omega}}(x(\hat{\theta},\hat{s},\kappa)) \mid \theta,s,\kappa] - T\} \qquad (28) \\ \leq \quad & \sum_{k:s-k\in\Omega,k\neq\kappa} Pr(k \mid \theta,s)\{E_{\tilde{\omega}}[v_{\theta\tilde{\omega}}(x(\hat{\theta},\hat{s},k)) \mid \theta,s,k] - t(\hat{\theta},\hat{s},k)\} \\ & + Pr(\kappa \mid \theta,s)E_{\tilde{\omega}}[v_{\theta\tilde{\omega}}(x(\hat{\theta},\hat{s},\kappa)) \mid \theta,s,\kappa] - bT. \qquad (29) \end{aligned}$$

This expression becomes smaller than $u(\theta,s;\hat{\theta},s)$ in (26) when T gets large. This shows that if T is sufficiently large, then after any ex ante report $\hat{\theta}$, agent type θ reports s truthfully under (x,t).

To see (ii), I compute $U_{\theta,\hat{\theta}}$. By (i), it is optimal for the agent to report s truthfully ex post for any ex ante report $\hat{\theta}$. Hence,

$$U_{\theta,\hat{\theta}} = \sum_{\omega,k,s} p_\theta(\omega)\mu^0(k)\pi^0_{\omega k}(s)[v_{\theta\omega}(x(\hat{\theta},s,k)) - t(\hat{\theta},s,k)]. \tag{30}$$

To understand the sum, fix ω and k and consider the summation over s. By definition of Π^0, $\pi^0_{\omega k}(s) = 1$ if $s = \omega + k$, and $\pi^0_{\omega k}(s) = 0$ for all $s \neq \omega + k$. Moreover, by definition of the mechanism, we have for $s = \omega + k$:

$$x(\hat{\theta},s,k) = x^*_{\hat{\theta},\omega}, \quad \textit{and} \quad t(\hat{\theta},s,k) = \tau(\hat{\theta},\omega). \tag{31}$$

Therefore, we obtain:

$$\sum_s p_\theta(\omega)\mu(k)\pi^0_{\omega k}(s)[v_{\theta\omega}(x(\hat{\theta},s,k)) - t(\hat{\theta},s,k)] = p_\theta(\omega)\mu^0(k)\cdot[v_{\theta\omega}(x^*_{\hat{\theta},\omega}) - \tau(\hat{\theta},\omega)]. \tag{32}$$

Hence, when now summing over ω and k, we obtain that

$$V_{\theta,\hat{\theta}} = \sum_\omega p_\theta(\omega)[v_{\theta\omega}(x^*_{\hat{\theta},\omega}) - \tau(\hat{\theta},\omega)]\left(\sum_k \mu^0(k)\right) = \sum_\omega p_\theta(\omega)[v_{\theta\omega}(x^*_{\hat{\theta},\omega}) - \tau(\hat{\theta},\omega)]. \tag{33}$$

By inspection, $U_{\theta,\hat{\theta}}$ coincides with the function that appears on the right hand side of (23), and $U_{\theta,\theta}$ coincides with the functions that appear on the left hand sides of (23) and (24). Since τ is a solution to $\tilde{P}$, it follows from (23) that $U_{\theta,\theta} \geq U_{\theta,\hat{\theta}}$ for all $\hat{\theta},\theta$, and hence the mechanism (x,t) is ex ante incentive compatible.

As to (iii), since (24) is binding under the solution τ to $\tilde{P}$, (33) implies that $U_{\theta,\theta} = 0$, and hence the mechanism (x,t) is individually rational.

As to (iv), analogous steps as in (ii) yield that the principal's utility from the mechanism (x,t) coincides with the objective (22) in $\tilde{P}$. Since the value of $\tilde{P}$ is equal to Z^*, and no agent type gets an information rent under (x,t) by (iii), the principal obtains Z^* from (x,t). qed

Proof of Proposition 2 The proof is analogous to the proof of Proposition 1. Therefore, I sketch only the main differences. Let information structure (16) be given. I say that $s_i = (\sigma_i, k_j)$ and $s_j = (\sigma_j, k_i)$ are consistent if

$$\sigma_i - k_i \in \Omega_i \quad \text{and} \quad \sigma_j - k_j \in \Omega_j, \quad i \neq j, \tag{34}$$

and are inconsistent otherwise. Lemma 1 is replaced by

Lemma A.1. *(i) There is* $\mu^0 \in \mu(K)$ *and* $\beta > 0$ *so that for all* θ_i, s_i *and all* s_j *that are consistent with* s_i*, we have* $Pr(s_j \mid \theta_i, s_i) > \beta$.
(ii) For all signals s_i *and all reports* $\hat{s}_i$*, there is a* s_j*,* $j \neq i$*, so that* s_i *and* s_j *are consistent, but* $\hat{s}_i$ *and* s_j *are inconsistent.*

To see part (i) of the proof Lemma A.1, observe that because of independence,

$$Pr(s_j \mid \theta_i, s_i) = Pr(\sigma_j \mid \theta_i, s_i) Pr(k_i \mid \theta_i, s_i) \tag{35}$$

The same steps as in Lemma 1, (i), imply that $Pr(k_i \mid \theta_i, s_i)$ is bounded from below. Moreover,

$$Pr(\sigma_j \mid \theta_i, s_i) = \sum_{\theta_j} p^{(j)}_{\theta_j}(\sigma_j - k_j) r_j(\theta_j) \geq \min_{\theta_j, \omega_j} p^{(j)}_{\theta_j}(\omega_j) \tag{36}$$

is bounded from below.

To see part (ii), let $s_i = (\sigma_i, k_j)$ and $\hat{s}_i = (\hat{\sigma}_i, \hat{k}_j)$ with $\hat{s}_i \neq s_i$ be given. If $\hat{\sigma}_i < \sigma_i$, then for $s_j = (\sigma_j, k_i)$ with $k_i = \sigma_i + 1$, we have that $\sigma_i - k_i = 1 \in \Omega_i$, but $\hat{\sigma}_i - k_i < 1 \notin \Omega_i$. If $\hat{\sigma}_i > \sigma_i$, then for $s_j = (\sigma_j, k_i)$ with $k_i = \sigma_i - \bar{\omega}_i$, we have that $\sigma_i - k_i = \bar{\omega}_i \in \Omega_i$, but $\hat{\sigma}_i - k_i > \bar{\omega}_i \notin \Omega_i$.

Moreover, if $\hat{k}_j < k_j$, then for $s_j = (\sigma_j, k_i)$ with $\sigma_j = k_j + \bar{\omega}_j$, we have that $\sigma_j - k_j = \bar{\omega}_j \in \Omega_j$, but $\sigma_j - \hat{k}_i > \bar{\omega}_j \notin \Omega_j$. If $\hat{k}_j > k_j$, then for $s_j = (\sigma_j, k_i)$ with $\sigma_j = k_j + 1$, we have that $\sigma_j - k_j = 1 \in \Omega_j$, but $\sigma_j - \hat{k}_i < 1 \notin \Omega_j$. qed

With the help of Lemma A.1, the construction of the mechanism in the single agent case can be extended to the multiple agents case as follows. For functions $\tau_i : \Theta_i \times \Omega_i \to \mathbb{R}$, define

$$\tilde{W} = \sum_{\theta \in \Theta, \omega \in \Omega} r_1(\theta_1) p^{(1)}_{\theta_1}(\omega_1) r_2(\theta_2) p^{(2)}_{\theta_2}(\omega_2) [\tau_1(\theta_1, \omega_1) + \tau_2(\theta_2, \omega_2) + v^P(x^*_{\theta\omega})], \tag{37}$$

$$\tilde{V}^{(i)}_{\theta_i, \hat{\theta}_i} = \sum_{\omega_i \in \Omega_i} p^{(i)}_{\theta_i}(\omega_i) \{ E[v^{(i)}_{\theta\omega}(x^*_{\hat{\theta}_i, \theta_j, \omega}) \mid \theta_i, \omega_i] - \tau_i(\hat{\theta}_i, \omega_i) \}. \tag{38}$$

The multi-agent version of the auxiliary problem $\tilde{P}$ in Proposition 1 is given by

$$\tilde{R}: \quad \max_{\tau_1(\theta_1, \omega_1), \tau_2(\theta_2, \omega_2)} \tilde{W} \quad s.t. \quad \tilde{V}^{(i)}_{\theta_i, \theta_i} \geq \tilde{V}^{(i)}_{\theta_i, \hat{\theta}_i}, \quad \tilde{V}^{(i)}_{\theta_i, \theta_i} \geq 0 \quad \forall i, \theta_i, \hat{\theta}_i.$$

Now, if the principal has full informational control, then a solution to problem $\tilde{R}$ is given by $\tau_i(\hat{\theta}_i, \omega_i) = E[v^{(i)}_{\theta\omega}(x^*_{\hat{\theta}_i, \theta_j, \omega}) \mid \theta_i, \omega_i]$, which is independent of θ_i by (13) and (14). Indeed, in this case, incentive compatibility $\tilde{V}^{(i)}_{\theta_i, \theta_i} \geq \tilde{V}^{(i)}_{\theta_i, \hat{\theta}_i}$, is trivially satisfied, and the individual rationality constraint is binding: $\tilde{V}^{(i)}_{\theta_i, \theta_i} = 0$ for $i = 1, 2$. If the principal has only partial informational control, then it follows again as in Riordan and Sappington (1988) that there is a solution $\tau_i(\theta_i, \omega_i)$ to

$\tilde{R}$ with $\tilde{V}^{(i)}_{\theta_i,\theta_i} = 0$ for $i = 1,2$, if the beliefs $p^{(i)}_{\theta_i}$ satisfy the spanning condition. In either case, the value of the problem is the complete information first-best surplus Z^*.

I now define a candidate solution (x, t_1, t_2) to the principal's original problem that will deliver Z^*. Recall that $s_i = (\sigma_i, k_j)$, $i \neq j$ with $\sigma_i = \omega_i + k_i$. Let

$$x(\theta, s_1, s_2) = \begin{cases} x^*_{\theta,\sigma_1-k_1,\sigma_2-k_2} & if \quad s_\ell - k_\ell \in \Omega_\ell, \ell = 1,2, \\ x^0 & else \end{cases} \tag{39}$$

$$t_i(\theta, s_1, s_2) = \begin{cases} \tau_i(\theta_i, \sigma_i - k_i) & if \quad s_\ell - k_\ell \in \Omega_\ell, \ell = 1,2, \\ T_i & else \end{cases} \tag{40}$$

for some $T_i > 0$.

Applying similar steps as in the proof of Proposition 1, it can be shown that, in the original problem, (x, t_1, t_2) is for each agent (i) ex post incentive compatibility, (ii) ex ante incentive compatible, (iii) individually rational, and (iv) yields the principal Z^*. qed

Proof of Proposition 3 Because the spanning condition fails, there are $\tilde{\theta}$ and $\alpha_\theta \in [0,1]$, $\theta \neq \tilde{\theta}$, with $\sum_{\theta\neq\tilde{\theta}} \alpha_\theta = 1$ so that

$$p_{\tilde{\theta}}(\omega) = \sum_{\theta\neq\tilde{\theta}} \alpha_\theta p_\theta(\omega) \quad \forall \omega. \tag{41}$$

Moreover, consider a valuation function with the property that for all θ, ω:

$$v_{\tilde{\theta}\omega}(x^*_{\tilde{\theta}\omega}) < v_{\theta\omega}(x^*_{\tilde{\theta}\omega}). \tag{42}$$

Let S and K be arbitrary measure spaces endowed with σ-algebras, and consider a compound information structure given by probability measures $\pi_{\omega k}(\cdot;\hat{\theta}) \in \Delta(S)$ and $\mu \in \Delta(K)$ capturing the (report-dependent) conditional signal distributions and the principal's randomization strategy. Towards a contradiction, suppose that the principal can extract the first-best surplus, then by the revelation principle, the agent reports truthfully ex ante and ex post on the equilibrium path, and in state ω, the first-best quantity $x^*_{\theta\omega}$ is implemented. Hence, type $\tilde{\theta}$'s utility is

$$U_{\tilde{\theta},\tilde{\theta}} = \sum_\omega \int_K \int_S p_{\tilde{\theta}}(\omega)[v_{\tilde{\theta}\omega}(x^*_{\tilde{\theta},\omega}) - t(\tilde{\theta}, s, k)]\, d\pi_{\omega k}(s;\tilde{\theta}) d\mu(k) \tag{43}$$

$$= \sum_{\theta\neq\tilde{\theta}} \alpha_\theta \left(\sum_\omega \int_K \int_S p_\theta(\omega)[v_{\tilde{\theta}\omega}(x^*_{\tilde{\theta},\omega}) - t(\tilde{\theta}, s, k)]\, d\pi_{\omega k}(s;\tilde{\theta}) d\mu(k) \right) \tag{44}$$

$$< \sum_{\theta\neq\tilde{\theta}} \alpha_\theta \left(\sum_\omega \int_K \int_S p_\theta(\omega)[v_{\theta\omega}(x^*_{\tilde{\theta},\omega}) - t(\tilde{\theta}, s, k)]\, d\pi_{\omega k}(s;\tilde{\theta}) d\mu(k) \right) \tag{45}$$

$$\leq \sum_{\theta\neq\tilde{\theta}} \alpha_\theta U_{\theta,\tilde{\theta}}, \tag{46}$$

where I have used (41) and (42) in the second and third line. To understand the final inequality, notice that the expression in the brackets in (45) is the utility of agent type θ when he (untruthfully) reports $\tilde{\theta}$ ex ante and reports s truthfully ex post. But, because after an untruthful report ex ante, it is not necessarily optimal to report truthfully ex post, this expression is (weakly) smaller than $U_{\theta,\tilde{\theta}}$ which, by definition, is agent type θ's utility when he (untruthfully) reports $\tilde{\theta}$ ex ante and chooses an optimal report ex post.

Now, because the principal extracts the full surplus, all agent types θ get $U_{\theta,\theta} = 0$. Together with incentive compatibility, this implies that $U_{\theta,\tilde{\theta}} \leq U_{\theta,\theta} = 0$, and hence the inequality above implies that $U_{\tilde{\theta},\tilde{\theta}} < 0$, contradicting individual rationality for type $\tilde{\theta}$. qed

Proof of Proposition 4 Let $S = K = \mathbb{Z}$, and take $\mu^0 \in \Delta(K)$ from Lemma 1. For some $s_0 \in S$, define

$$\pi_{\omega k}(s;\theta,\hat{\theta}) = \begin{cases} 1 & if \quad \theta = \hat{\theta} \text{ and } s - k = \omega, \\ 1 & if \quad \theta \neq \hat{\theta} \text{ and } s = s_0, \\ 0 & \quad\ else. \end{cases} \tag{47}$$

Hence, if the agent reports his type truthfully, the information structure coincides with (9), and if he misrepresents his type, he gets the signal s_0 with probability 1 which is therefore entirely uninformative, both about ω and k. Define the mechanism as follows:

$$x(\theta,s,k) = \begin{cases} x^*_{\theta,s-k} & if \quad s-k \in \Omega \\ 0 & if \quad s-k \notin \Omega \end{cases}, \quad t(\theta,s,k) = \begin{cases} v_{\theta,s-k}(x^*_{\theta,s-k}) & if \quad s-k \in \Omega \\ T & if \quad s-k \notin \Omega \end{cases} \tag{48}$$

for some $T > 0$.

For sufficiently large T, it follows as in the proof of Proposition 1 that the agent reports s truthfully, if he has reported θ truthfully ex ante. Therefore, the definition of payments, and the fact that k and s reveal the true state, implies that the agent obtains utility 0 when he reports θ truthfully.

Next, consider the case that agent type θ falsely reports $\hat{\theta} \neq \theta$ ex ante. Then the agent observes s_0 for sure and chooses an optimal report $\hat{s}(\theta)$ ex post. Because

$$\hat{s}(\theta) - k \in \Omega \quad \Leftrightarrow \quad k \in \{\hat{s}(\theta) - \bar{\omega}, \ldots, \hat{s}(\theta) - 1\}, \tag{49}$$

the agent, at the ex ante reporting stage, anticipates that with (at least) probability $\beta = \min_\theta \sum_{\ell \notin \{\hat{s}(\theta)-\bar{\omega},\ldots,\hat{s}(\theta)-1\}} \mu^0(\ell) > 0$, he receives a quantity of 0 and has to make payments T. Since β is independent of θ, it follows that for sufficiently large T, any agent type θ's expected utility from lying ex ante becomes negative. Because the agent's utility from truth-telling is 0, he is deterred from lying.

Moreover, because the mechanism implements the first-best quantities and the agent receives 0 rent, the principal fully extracts the first-best surplus Z^*. qed

B Appendix

The objective of this appendix is to clarify the relation between my and Zhu's (2018) construction for the case with two agents and quasi-linear utility. To illustrate, suppose there are two states per agent $\omega_i \in \{1,2\}$, and thus four states $\omega \in \{1,2\}^2$. Zhu (2018) considers an information structure with five signals $s_i \in \{1,\ldots,5\}$ per agent, illustrated in Table 3. A column depicts the pairs of signals $s = (s_1, s_2)$ that are jointly released to agents with probability $1/5$ each, conditional on the state ω depicted in the top row.

$s = (s_1, s_2)$	$\omega = (1,1)$	$\omega = (1,2)$	$\omega = (2,1)$	$\omega = (2,2)$
	$(1,2)$	$(1,3)$	$(1,4)$	$(1,5)$
	$(2,3)$	$(2,4)$	$(2,5)$	$(2,1)$
	$(3,4)$	$(3,5)$	$(3,1)$	$(3,2)$
	$(4,5)$	$(4,1)$	$(4,2)$	$(4,3)$
	$(5,1)$	$(5,2)$	$(5,3)$	$(5,4)$

Table 3: Zhu's (2018) information structure

Note the two key features: as each individual signal s_i is equally likely to occur in any state, agent i's belief about the state is unaffected by s_i. Individual signals are thus entirely uninformative about the state. But, since a pair (s_1, s_2) occurs in exactly one state, jointly s_1 and s_2 reveal ω.

Zhu (2018) shows that, when agents have to report their types and signals simultaneously, the designer can implement the same outcome as in the benchmark in which agents only know their types and the state becomes public ex post. Zhu (2018) shows first that signals can be elicited at no cost, thus revealing the state. Intuitively then, since individual signals are uninformative about the state, their presence does not make it harder to meet incentive constraints (with respect to types) than in the benchmark.

To elicit signals at no cost, Zhu (2018) exploits that a signal s_i is informative about the other agent's signal s_j, $j \neq i$. Hence, agents' posterior beliefs about the other agent's signal are correlated. Specifically, agent 1's beliefs about agent 2 are summarized by the matrix $B_1(\theta_1)$ below. The entry in the m-th row and the n-th column depicts the probability $Pr(s_2 = n \mid \theta_1, s_1 = m)$ which agent 1 attaches to the event that agent 2 has observed $s_2 = n$, conditional on agent 1 having observed θ_1 and $s_1 = m$. It is not hard to see from Table 1 that these beliefs correspond to agent 1's belief $Pr((\omega_1, \omega_2) \mid \theta_1)$ that the state is (ω_1, ω_2) after having observed θ_1.

$$B_1(\theta_1) = \begin{pmatrix} 0 & Pr((1,1)\mid\theta_1) & Pr((1,2)\mid\theta_1) & Pr((2,1)\mid\theta_1) & Pr((2,2)\mid\theta_1) \\ Pr((2,2)\mid\theta_1) & 0 & Pr((1,1)\mid\theta_1) & Pr((1,2)\mid\theta_1) & Pr((2,1)\mid\theta_1) \\ Pr((2,1)\mid\theta_1) & Pr((2,2)\mid\theta_1) & 0 & Pr((1,1)\mid\theta_1) & Pr((1,2)\mid\theta_1) \\ Pr((1,2)\mid\theta_1) & Pr((2,1)\mid\theta_1) & Pr((2,2)\mid\theta_1) & 0 & Pr((1,1)\mid\theta_1) \\ Pr((1,1)\mid\theta_1) & Pr((1,2)\mid\theta_1) & Pr((2,1)\mid\theta_1) & Pr((2,2)\mid\theta_1) & 0 \end{pmatrix}. \tag{50}$$

If the type θ_1 is orthogonal to the state ω, then $B_1(\theta_1) = B_1$ is independent of θ_1, and B_1 has typically full rank. Zhu (2018) uses these properties to constructs payments as in Cremer and McLean (1988) that elicit signals (s_1, s_2) at no cost.

Observe that the information structure in Table 3 does satisfy my constraint C. Thus Zhu's (2018) result, a fortiori, carries over to my setting when agents report type and signal sequentially, and constraint C is imposed. Unlike me, Zhu (2018) does not, however, discuss the case in which θ_1 is not orthogonal to ω_1. In this case, $B_1(\theta_1)$ may still have full rank, but since θ_1 is the agent's private information, the payments to elicit signals would need to condition on a report about θ_1. It is an open question whether (even if agents report sequentially), using the information structure of Table 3, payments can be constructed that elicit θ_i and s_i at no cost.

References

Baron, D. and D. Besanko (1984). Regulation and information in a continuing relationship. Information Economics and Policy, 1, 267-302.

Bergemann, D., A. Bonatti, and A. Smolin (2017). The design and price of information. American Economic Review, forthcoming.

Bergemann, D. and S. Morris (2016). Bayes correlated equilibrium and the comparison of information structures in games. Theoretical Economics, 11, 487-522.

Bergemann, D. and M. Pesendorfer (2007). Information structures in optimal auctions. Journal of Economic Theory, 137, 580-609.

Bergemann, D. and A. Wambach (2015). Sequential information disclosure in auctions. Journal of Economic Theory, 159, 1074-1095.

Courty, P. and H. Li (2000). Sequential screening. Review of Economic Studies, 67, 697-717.

Crémer, J. and R. P. McLean (1988). Full extraction of the surplus in Bayesian and dominant strategy auctions. Econometrica, 56, 1247-1257.

Esö, P. and B. Szentes (2007a). Optimal information disclosure in auctions and the handicap auction. Review of Economic Studies, 74, 705-731.

Esö, P. and B. Szentes (2007b). The price of advice. The Rand Journal of Economics, 38, 863-880.

Esö, P. and B. Szentes (2017). Dynamic contracting: an irrelevance result. Theoretical Economics, 12, 109-139.

Kamenica, E. and M. Gentzkow (2011). Bayesian persuasion. American Economic Review, 101, 2590-2615.

Kolotilin, A., M. Li, T. Myolovanov, and A. Zapechelnyuk (2017). Econometrica, 85, 1949-1964.

Krähmer, D. (2012). Optimal auction design with endogenously correlated buyer types. Journal of Economic Theory, 147, 118-141.

Krähmer, D. and R. Strausz (2015a). Ex post information rents in sequential screening. Games and Economic Behavior, 90, 257-273.

Krähmer, D. and R. Strausz (2015b). "Dynamic mechanism design." Chapter 11 in ed. T. Börgers, An Introduction to the Theory of Mechanism Design, 204-234.

Myerson, R. B. (1982). Optimal coordination mechanisms in generalized principal-agent problems. Journal of Mathematical Economics, 10, 67-81.

Myerson, R. (1986). Multistage games with communication. Econometrica, 54, 323-358.

Li, H. and Shi, X. (2017a). Discriminatory information disclosure. American Economic Review, 107, 3363-3385.

Li, H. and Shi, X. (2017b). Optimal discriminatory disclosure. Working paper, mimeo.

McAfee, P. and P. Reny (1992). Correlated information and mechanism design. Econometrica, 60, 395-421.

Obara, I. (2008). The full surplus extraction theorem with hidden actions. The B.E. Journal of Theoretical Economics, 8 (Advances), Article 8.

Rahman, D. (2012). But who will monitor the monitor?. The American Economic Review, 102, 2767-2797.

Rahman, D. and I. Obara (2010). Mediated partnerships. Econometrica, 78, 285-308.

Riordan, M. H. and D. E. Sappington (1988). Optimal contracts with public ex post information. Journal of Economic Theory, 45, 189-199.

Rodina, D. (2018). Information design and career concerns. Working paper, Northwestern University, mimeo.

Smolin, A. (2018). Disclosure and pricing of attributes. Working paper, University Bonn, mimeo.

Strausz, R. (2012). Mediated contracts and mechanism design. Journal of Economic Theory, 147, 1280-1290.

Zhu, S. (2018). Private disclosure with multiple agents. Working paper, Toulouse School of Economics, mimeo.