

Intervention with Private Information, Imperfect Monitoring and Costly Communication

Luca Canzian, *Member, IEEE*, Yuanzhang Xiao, William Zame, Michele Zorzi, *Fellow, IEEE*,
and Mihaela van der Schaar, *Fellow, IEEE*

Abstract—This paper studies the interaction between a designer and a group of strategic and self-interested users who possess information the designer does not have. Because the users are strategic and self-interested, they will act to their own advantage, which will often be different from the interest of the designer, even if the latter is benevolent and seeks to maximize (some measure of) social welfare. In the settings we consider, the designer and the users can communicate (perhaps with noise), the designer can observe the actions of the users (perhaps with error) and the designer can commit to (plans of) actions – *interventions* – of its own. The designer’s problem is to construct and implement a *mechanism* that provides incentives for the users to communicate and act in such a way as to further the interest of the designer – *despite* the fact that they are strategic and self-interested and possess private information. To address the designer’s problem we propose a general and flexible framework that applies to many scenarios. To illustrate the usefulness of this framework, we discuss some simple examples, leaving further applications to other papers. In an important class of environments, we find conditions under which the designer can obtain its benchmark optimum – the utility that could be obtained if it had all information and could command the actions of the users – and conditions under which it cannot. More broadly we are able to characterize the solution to the designer’s problem, even when it does not yield the benchmark optimum. Because the optimal mechanism may be difficult to construct and implement, we also propose a simpler and more readily implemented mechanism that, while falling short of the optimum, still yields the designer a “good” result.

Index Terms—Game theory, incomplete information, mechanism design, intervention, resource allocation.

I. INTRODUCTION

WE study the interaction between a group of users and a designer. If the users are compliant or the designer can command the actions of the users, then the designer is faced with an optimal control problem of the sort that is well-studied. Little changes if the users have private information (about themselves or about the environment) that the designer does not have but the designer can communicate with the users, because the designer can simply ask or instruct the

users to report that information. However, a great deal changes if the users are not compliant but rather are self-interested and strategic and the designer can not command the actions and reports of the users. In that case, the users may (and usually will) take actions and/or provide reports that are in their own self-interest but not necessarily in the interest of the designer. The objective of this work is to understand the extent to which the designer can provide incentives to the users to take actions and provide reports that further the objectives of the designer, be those selfish or benevolent. (The case of a benevolent designer is probably the one of most interest, but the problem faced by a benevolent designer is no easier than the problem faced by a selfish designer: the goal of a benevolent designer is to maximize some measure of social welfare – which might include both total utility and some measure of fairness – but the goal of an individual user is to maximize its own utility; hence the incentives of the designer and of the individual user are no more aligned when the designer is benevolent than when the designer is selfish, so the same incentives to misrepresent and misbehave are present in both circumstances. Such incentives frequently lead to the over-use of resources and to substantial inefficiencies [1], [2].) Here, we are specifically interested in settings in which the users can send reports to the designer and the designer in turn can send messages to the users before the users act, after which the designer may take actions of its own – *interventions*. Our use of intervention builds on [3]–[6], but we go beyond that work in considering private information, imperfect monitoring and costly communication – in addition to intervention itself.

Our work has something in common with the economic theory of *mechanism design* in the tradition of [7]–[11]. Indeed, our general framework builds on that of [12], and the abstract theory of mechanism design – in particular the revelation principle – does play a role. However, [12] does not solve any of our problems because after we use the revelation principle to restrict our attention to incentive compatible direct mechanisms we must still *construct* the optimal mechanism, which is a non-trivial undertaking.¹ Moreover, when we admit physical (and other) constraints, noisy communication and imperfect monitoring, the revelation principle does not help because it entirely obscures all of these complications. Finally, the revelation principle simply does not hold when communication is costly.

Manuscript received August 3, 2012; revised January 7 and April 18, 2013. The editor coordinating the review of this paper and approving it for publication was A. MacKenzie.

L. Canzian, Y. Xiao, and M. van der Schaar are with the Department of Electrical Engineering, UCLA, Los Angeles CA 90095, USA (e-mail: luca.canzian@gmail.com).

W. Zame is with the Department of Economics, UCLA, Los Angeles CA 90095, USA.

M. Zorzi is with the Department of Information Engineering, University of Padova, Via Gradenigo 6/B, 35131 Padova, Italy.

The work of Luca Canzian, Yuanzhang Xiao and Mihaela van der Schaar was partially supported by the NSF grant CCF 1218136.

The work of Michele Zorzi was partially supported by Fondazione Cariparo through the program “Progetti di Eccellenza 2011-2012.”

Digital Object Identifier 10.1109/TCOMM.2013.061013.120558

¹Proposition 1 in [12] shows that the problem of choosing the optimal incentive compatible direct mechanism is a linear programming problem *provided* that type sets and action sets are finite *and fixed* and that the designer can send arbitrary messages – but in our context the action sets may not be finite, the action set of the designer is definitely not fixed, and the designer’s choice of messages may be constrained, so our problem is different, and much more complicated.

We treat settings in which the users have private information but (perhaps limited and imperfect) communication between the users and the designer – more precisely, the *device* employed by the designer – is possible. The users have the opportunity to send reports about their private information, and the device can in turn send messages to the users; in both cases, we allow for the possibility that communication is noisy so that the report/message sent is not necessarily the report/message received. After this exchange of information, the users take actions. Finally, the device, having (perhaps imperfectly) observed these actions, also has the opportunity to act. Generalizing a construction of [12], we formalize this setting as a *communication mechanism*. The device the designer employs plays two roles: first, to *coordinate* the actions of the users *before* they take them, and second, to *discipline* the users *afterwards*. Because users are self-interested and strategic, their reports and actions will only serve the interest of the designer if they also serve their own interests. Thus we are interested in strategy profiles for the users that each user finds optimal, given the available information, the strategies of others and the nature of the given device; we refer to these as *communication equilibria*. Note that the device is not strategic – it is a device after all – but *the designer behaves strategically in choosing the device*. Because we focus here on the problem of the designer, we are interested in finding devices that support equilibria that the designer finds optimal. (If the designer is benevolent – i.e., intends to maximize social welfare, perhaps constrained by some notion of fairness – these devices will also serve the interests of the users as a whole, but if the designer is self-interested they may not.) We are particularly interested in knowing when the designer can find a device so as to achieve his benchmark optimum – the outcome he could achieve if he knew all relevant information and users were fully compliant – despite the fact that information is in fact private and users are in fact self-interested. For a class of environments that includes many engineering problems of interest (e.g., power control [6], [13], medium access control (MAC) [4], [14], and flow control [14]–[18]) we find conditions under which there exist mechanisms that achieve the benchmark optimum and conditions under which such mechanisms do not exist. In case they do not exist, we find conditions such that the problem of finding an optimal protocol can be decoupled. Because the optimal protocol may still be difficult to compute, we also provide a simple algorithm that converges to a protocol that, although perhaps not optimal, still yields a “good” outcome for the designer. [19] demonstrates that these non-optimal protocols are very useful in a flow control management system.

Throughout, we assume that the designer can *commit* to a choice of a device that is pre-programmed to carry out a particular plan of action *after* the reports and actions of the users. In mechanical terms, such commitment is possible precisely because the designer deploys a device – hardware or software or both – *and then leaves*. Indeed, the desire of the designer to commit is one reason that it employs a device. Although other assumptions are possible, this assumption seems most appropriate for the settings we have in mind, in which the designer is a long-lived and experienced entity who has learned the relevant parameters (user utilities and distribution

of user types) over time, but the users are short-lived, come and go but do not interact repeatedly: in a particular session they are not playing a repeated game and are not forward-looking.²

There is by now a substantial communication engineering literature that addresses the problem of providing incentives for strategic users to obey a particular resource allocation scheme. Such incentives might be provided in a number of different ways. [3]–[6] provide incentives via intervention, and the frameworks adopted in [15] and [20] capture the concept of intervention as well. However, none of the preceding works has considered the strategic aspect associated to the information acquisition, which is the qualifying point of our analysis. In this work we extend the intervention framework to deal with situations in which the users have private information and try to exploit such advantage, and we propose a novel methodology to retrieve the private information providing the users with an incentive to declare this information truthfully. A rather different literature, including [21]–[24], adopts literal pricing schemes in which users are required to make monetary payments for resource usage.³ Literal pricing schemes require the designer to have specific knowledge (the value of money to the users) and require a technology for making monetary transfers, which is missing in many settings, such as wireless communication. Moreover, it does not necessarily solve the problem of a benevolent designer since monetary payments are by definition costly for the user making them and hence wasteful.⁴ An additional difficulty in employing literal payment schemes is that it is debatable whether users would agree to a pricing scheme that dynamically varies with the state of the system, in particular if users have to pay for a service that had hitherto been free. A smaller literature [27]–[29] addresses environments in which users have private information – their private monetary valuations for access to the resource – and uses ideas from mechanism design and auction theory [30] to create protocols in which users are asked to report their private monetary valuations, after which access to the resource is apportioned and users make monetary payments according to their access and the reports of valuations. For very detailed comparison of intervention, pricing, and other approaches, see [31].

The remainder of this paper is organized as follows. Section II introduces our framework of devices and mechanisms and the notion of equilibrium. Section III presents an example to illustrate how private information, information revelation and intervention all matter. Section IV asks when some devices achieve the benchmark optimum. Section V studies the properties of the optimal devices and Section VI offers a

²If the interaction between the designer and the users is long and sustained, then the interaction should be modeled as a repeated game, in which case the analysis would be completely different and the designer should consider only policies which are part of a subgame perfect equilibrium.

³A different line of work, which includes [13], [25], [26] but is quite far from the work here, uses pricing in scenarios where users are compliant, rather than self-interested and strategic. In those scenarios, however, the function of pricing is decentralization: prices induce utility functions for the users that lead them to take the desired actions without the need for centralized control. In these scenarios pricing is figurative rather than literal: monetary payments are not actually required.

⁴A possible exception occurs when/if all users value money in the same way and it is possible to make transfers among users, so that some users make payments and others receive payments.

constructive procedure for choosing devices that are simple to compute and implement – if not necessarily optimal. Section VII concludes with some remarks.

II. FRAMEWORK

We consider a *designer* and a collection of *users*. The designer chooses an *intervention device* and then leaves – the designer itself takes no further actions. In a single *session* the device interacts with a fixed number of users n , labeled from 1 to n . We will write $N = \{1, \dots, n\}$ for the set of users. We think of the users in a particular session as drawn from a pool of potential users, so users may be (and typically will be) different in each session. We allow for the possibility that users are drawn from different pools – e.g., occupy different geographical locations or utilize different channels.

User i is characterized by an element of a set T_i of *types*; a user's type encodes all relevant information about the user, which will include the user's utility function and the influence the user's type has on other users and on the designer. Write $T = T_1 \times \dots \times T_n$ for the set of possible type profiles. Users know their own type; users and the designer know the distribution of user types π (a probability distribution on T).⁵ If user i is of type t_i then $\pi(\cdot | t_i)$ is the conditional distribution of types of other users. (We allow for the possibility that types are correlated, which might be the case, for instance, if users have private information about the current state of the world and not only about themselves.) In each session, user i chooses an *action* from the set A_i of *actions*. We write $A = A_1 \times \dots \times A_n$ for the set of possible action profiles and (a_i, a_{-i}) for the action profile in which user i chooses action $a_i \in A_i$ and other users choose the action profile $a_{-i} \in A_{-i} = A_1 \times \dots \times A_{i-1} \times A_{i+1} \times \dots \times A_n$; we use similar notation for types, etc.

The designer is characterized by its utility function and the set \mathcal{D} of *devices* it might use. A device – which might consist of hardware or software or both – has four features: 1) it can receive communications from users, 2) it can send communications to users, 3) it can observe the actions of users, and 4) it can take actions of its own. As in [3], we interpret the actions of the device as *interventions*. We formalize a device as a tuple $D = \langle (R_i), (M_i), \mu, X, \Phi, \epsilon^R, \epsilon^M, \epsilon^A \rangle$, where:

- R_i is the set of *reports* that user i might send; write $R = R_1 \times \dots \times R_n$ for the set of report profiles;
- M_i is the set of *messages* that the device might send to user i ; write $M = M_1 \times \dots \times M_n$ for the set of message profiles;
- $\mu : R \rightarrow \Delta(M)$ is the *message rule*, which specifies the (perhaps random) profile of messages to be sent to the users as a function of the reports received from all users; if r is the profile of observed reports we write μ_r for the corresponding probability distribution on M , and $\mu_r(m)$ for the probability that the message m is chosen when the observed report is r ;^{6,7}

⁵We usually think of the distribution π as common knowledge but this is not entirely necessary.

⁶ $\Delta(B)$, where B is a set with cardinality $b+1$, denotes the b -unit simplex.

⁷If the message m is always chosen given the observed report profile r , μ_r is point mass at m , i.e., $\mu_r(z) = 1$ if $z = m$, $\mu_r(z) = 0$ otherwise. However, in this case we usually prefer to abuse notation and write $\mu(r) = m$. Below, we will make similar notational abuses without further comment.

- X is the set of *interventions* (actions) the device might take;
- $\Phi : R \times M \times A \rightarrow \Delta(X)$ is the *intervention rule*, which specifies the (perhaps random) intervention the device will take given the received reports, the transmitted messages and the observed actions; if r are the observed reports, m the transmitted messages and a the observed actions, we write $\Phi_{r,m,a}$ for the corresponding probability distribution on X ;
- $\epsilon^R : R \rightarrow \Delta(R)$ encodes the noise in receiving reports: users send the report profile r but the designer observes a random profile \hat{r} distributed according to ϵ_r^R ;
- $\epsilon^M : M \rightarrow \Delta(M)$ encodes the noise in receiving messages: the device sends the message profile m but users observe a random profile \hat{m} distributed according to ϵ_m^M ;
- $\epsilon^A : A \rightarrow \Delta(A)$ encodes the error in monitoring actions: the users choose an action profile a but the device observes a random profile \hat{a} distributed according to ϵ_a^A .

The set of all conceivable devices is very large, but in practice the designer will need to choose a device from some prescribed (perhaps small) subset \mathcal{D} , so we assume this throughout. In this generality, reports and messages could be entirely arbitrary but typically reports will provide (perhaps incomplete) information about types, and messages will provide (perhaps incomplete) recommendations for actions, and we will frequently use this language.

If the report spaces are singletons then reports are meaningless, so singleton report spaces R_i express the absence of reporting. Similarly, a singleton message space M expresses the absence of messaging and a singleton intervention space X expresses the absence of intervention. The absence of noise/error with regard to reports, messages or actions can be expressed by requiring that the corresponding mapping(s) be the identity; e.g., ϵ_r^R is point mass at r and so $\hat{r} = r$ for all report profiles, etc. However, in any of these cases we would usually prefer to abuse notation and omit the corresponding component of the tuple that describes the device.

The utility $U_i(a, t, r_i, x)$ of user i depends on the actions a and types t of all users, the report r_i chosen by user i , and the intervention x of the designer. The utility $U(a, t, r, m, x)$ of the designer depends on the actions a and types t of all users, on the reports r , the messages m and the intervention x of the designer. The dependence of utility on reports and messages allows for the fact that communication may be costly. Note that the utility of a user depends only on the report that user sends, but the utility of the designer depends on the messages it sends *and* on the reports of the users. If this seems strange, keep in mind that if the designer is benevolent and seeks to maximize social utility, he certainly cares about the reporting costs of users.

A *communication mechanism*, or *mechanism* for short, is a tuple $\mathcal{C} = \langle N, (T_i, A_i, U_i), \pi, U, D \rangle$ that specifies the set N of users, the sets T_i of user types, the sets A_i of user actions, the utility functions U_i of users, the distribution π of types, the utility function U of the designer, and the device D . We view the designer as choosing the device, which is pre-programmed, but otherwise taking no part: the users choose and execute plans and the device carries out its programming.

The operation of a communication mechanism \mathcal{C} is as

follows.

- users make reports to the device;
- the device “reads” the reports (perhaps with error) and sends messages to the users (perhaps depending on the realization of the random rule);
- users “read” the messages⁸ (perhaps with error) and take actions;
- the device “monitors” the actions of the users (perhaps imperfectly) and, following the rule, makes an intervention (perhaps depending on the realization of the random rule).

A *strategy* for user i is a pair of functions $f_i : T_i \rightarrow R_i$, $g_i : T_i \times M_i \rightarrow A_i$ that specify which report to make, conditional on the type of user i , and which action to take, conditional on the type of user i and the message observed. We do not specify a strategy for the device because the device is not strategic, its behavior is completely specified by the message rule and the intervention rule – but the *designer behaves strategically in choosing the device*. Given a profile (f, g) of user strategies, and the intervention device D , the expected utility of a user i whose type is t_i is obtained by averaging over all random variables involved, i.e.,⁹:

$$EU_i(f, g, t_i, D) = \sum_{t_{-i} \in T_{-i}} \pi(t | t_i) \sum_{\hat{r} \in R} \epsilon_r^R(\hat{r}) \sum_{m \in M} \mu_{\hat{r}}(m) \cdot \sum_{\hat{m} \in M} \epsilon_m^M(\hat{m}) \sum_{\hat{a} \in A} \epsilon_a^A(\hat{a}) \sum_{x \in X} \Phi_{\hat{r}, m, \hat{a}}(x) U_i(a, t, r_i, x)$$

where $r_j = f_j(t_j)$ and $a_j = g_j(t_j, \hat{m}_j)$, $\forall j \in N$, are the reports sent and the actions taken by users. Similarly, the expected utility of the designer is

$$EU(f, g, D) = \sum_{t \in T} \pi(t) \sum_{\hat{r} \in R} \epsilon_r^R(\hat{r}) \sum_{m \in M} \mu_{\hat{r}}(m) \cdot \sum_{\hat{m} \in M} \epsilon_m^M(\hat{m}) \sum_{\hat{a} \in A} \epsilon_a^A(\hat{a}) \sum_{x \in X} \Phi_{\hat{r}, m, \hat{a}}(x) U(a, t, r, m, x)$$

The strategy profile (f, g) is an *equilibrium* if each user is optimizing given the strategies of other users and the device D ; that is, for each user i we have

$$EU_i(f_i, f_{-i}, g_i, g_{-i}, t_i, D) \geq EU_i(f'_i, f_{-i}, g'_i, g_{-i}, t_i, D) \quad (1)$$

for all strategies $f'_i : T_i \rightarrow R_i$, $g'_i : T_i \times M_i \rightarrow A_i$.¹⁰ We often say that the device D *sustains* the profile (f, g) . We remark that the existence of such an equilibrium is not always guaranteed without additional assumptions and needs to be explicitly addressed in the specific case at hand.

The designer seeks to optimize his own utility by choosing a device D from some prescribed class \mathcal{D} of physically feasible

⁸Note that we assume that each user i can only read its own message m_i . However, our framework is suitable to model also situations in which user i is able to hear the message m_j intended for user j . In this case it is sufficient to focus on devices in which the message sent to user j is part of the message sent to user i .

⁹We have tacitly assumed that all the probability distributions under consideration have finite or countably infinite support – which will certainly be the case if the spaces under consideration are themselves finite or countably infinite; in a more general context we would need to replace summations by integrals and to be careful about measurability.

¹⁰The notion of equilibrium defined here is that of a Bayesian Nash equilibrium of the Bayesian game induced by the communication mechanism. For simplicity, we have restricted our attention to equilibrium in *pure strategies*; we could also allow for equilibria in *mixed strategies* [32].

devices. Because users are strategic, the designer must assume that, whatever device D is chosen, the users will follow some equilibrium strategy profile (f, g) . Since the designer will typically recommend actions, we assume that, if more than one equilibrium strategy profile exists, the users choose (because they are coordinated to) the equilibrium that the designer most prefers (in case of a benevolent manager, it usually coincides with the equilibrium that the users prefer). Hence, the designer has to solve the following **Optimal Device (OD)** problem¹¹:

$$\operatorname{argmax}_{D \in \mathcal{D}} \max_{f, g} EU(f, g, D)$$

subject to:

$$EU_i(f_i, f_{-i}, g_i, g_{-i}, t_i, D) \geq EU_i(f'_i, f_{-i}, g'_i, g_{-i}, t_i, D) \\ \forall i \in N, \forall t_i \in T_i, \forall f'_i : T_i \rightarrow R_i, \forall g'_i : T_i \times M_i \rightarrow A_i$$

We say that a solution D of the above problem is an *optimal device*. To maintain parallelism with some other literature, we sometimes abuse language and refer to the designer’s problem as choosing an *optimal mechanism* – even though the designer only chooses the device and not the types of users, their utilities, etc. Note that optimality is relative to the prescribed set \mathcal{D} of considered devices. Moreover, the expected utility the designer obtains choosing the optimal device must not be confused with the *benchmark optimum* utility the designer could achieve if users were compliant, which is in general higher. If they coincide, we say that the device D is a *maximum efficiency device*.

A. Null reports, messages and interventions

In many (perhaps most) concrete settings, it is natural to presume that users might sometimes choose not to make reports and that the device might sometimes not send messages or make an intervention. The easiest way to allow for these possibilities is simply to assume the existence of null reports, null messages and null actions. In particular, we can assume that for each user i there is a distinguished report r_i^* which is to be interpreted as “not sending a report”. (On the device side, observing r_i^* should be interpreted as “not receiving a report”.) Because not making a report should be costless, we should assume that, fixing types, reports of others to the device, actions by the users and intervention by the device, r_i^* yields utility at least as great as any other report: $U_i(a, t, r_i^*, x) \geq U_i(a, t, r_i, x)$ and $U(a, t, r_i^*, r_{-i}, m, x) \geq U(a, t, r_i, r_{-i}, m, x)$, for all $i, a, t, x, r_i, r_{-i}, m$. Given this assumption, and using utility when sending the report r_i^* as the baseline, we can interpret the differences $U_i(a, t, r_i^*, x) - U_i(a, t, r_i, x)$ and $U(a, t, r_i^*, r_{-i}, m, x) - U(a, t, r_i, r_{-i}, m, x)$ as the cost to user i and to the designer, respectively, of sending the report r_i . In this generality, the cost of sending a report might depend on all other variables. We remark that this cost does not take into consideration the impact of the communication on the interaction among the users and the intervention device: in deciding whether or not to send a report, a user must take into account the fact that

¹¹Because the utility functions of the users depend on reports, and the utility function of the designer depends on messages and reports, which are parameters of the device chosen, this tacitly assumes that utility functions are defined on a domain sufficiently large to encompass all the possibilities that may arise when any device $D \in \mathcal{D}$ is chosen.

sending a report may alter the messages sent by the device and hence the actions of the users and the intervention of the device. So sending a report, while increasing the instantaneous communication cost, may still lead to a higher utility because it influences the strategic choices of others.

Similarly, we could assume that for each user i there is a distinguished message m_i^* that the device might send but which we interpret as “not sending a message”. (On the user side, we interpret receipt of the message m_i^* as “not receiving a message”.) Because not sending a message m_i^* should be costless, we assume that $U(a, t, r, m_i^*, m_{-i}, x) \geq U(a, t, r, m_i, m_{-i}, x)$ for all $i, a, t, r, m_{-i}, x, m_i$, and so interpret the difference $U(a, t, r, m_i^*, m_{-i}, x) - U(a, t, r, m_i, m_{-i}, x)$ as the cost of sending the message m_i , which might depend on all other variables.

Finally, we could assume that there is a distinguished intervention x^* that we interpret as “not making an intervention”. If (as we usually do) we want to interpret an intervention as a *punishment*, we should assume that x^* yields utility at least as great as any other intervention for each user and the designer: $U_i(a, t, r_i, x^*) \geq U_i(a, t, r_i, x)$ and $U(a, t, r, m, x^*) \geq U(a, t, r, m, x)$ for all i, a, t, r, m, x , and we interpret the differences $U_i(a, t, r_i, x^*) - U_i(a, t, r_i, x)$ and $U(a, t, r, m, x^*) - U(a, t, r, m, x)$ as the cost of the intervention to user i and to the designer, respectively, which might depend on all other variables.

If the sets of reports (respectively, messages, interventions) are singletons, then by default there are no possible reports (respectively, messages, interventions).

If D is a device for which “not making an intervention” is possible and (f, g) is an equilibrium with the property that $\Phi_{\hat{r}, m, \hat{a}}(x^*) = 1$, for all type profiles t , observed reports \hat{r} , sent messages m and observed actions \hat{a} (with \hat{r} , m and \hat{a} occurring with positive probability), we say that D sustains (f, g) *without intervention*. The most straightforward interpretation is that the device threatens punishments for deviating from the recommended actions and that the threats are sufficiently severe that they do not need to be executed. Again, this is natural in context: by using the intervention device, the designer *commits* to meting out punishments for deviation, even if those punishments are costly for the designer as well as for the users.

B. Direct mechanisms

To be consistent with [12], we say that the mechanism C^d is a *direct mechanism* if $R_i = T_i$ for all i (users report their types, not necessarily truthfully) and $M = A$ (the device recommends action profiles), there are no errors, and reports and messages are costless (i.e., utility does not depend on reports or messages). If C^d is a direct mechanism we write $(f^*, g^*) = (f_1^*, \dots, f_n^*, g_1^*, \dots, g_n^*)$ for the strategy profile in which users are honest (report their true types) and obedient (follow the recommendations of the device); that is, $f_i^*(t_i) = t_i$ and $g_i^*(t_i, a_i) = a_i$ for every user i , type t_i , and recommendation $a_i = \mu_i(t_i)$. If (f^*, g^*) is an equilibrium, we say that C^d is *incentive compatible*. If a device is such that the resulting mechanism is an incentive compatible direct mechanism, we say that the device is incentive compatible.

Incentive compatible direct mechanisms play a special role because of the following generalization of the revelation principle. (We omit the proof, which is almost identical to the proof of Proposition 2 in [12].)

Proposition 1. *If C is a mechanism for which reports and messages are costless and (f, g) is an equilibrium of the mechanism C , then there is an incentive compatible direct mechanism C^d with the same action and intervention spaces for which the honest and obedient strategy profile (f^*, g^*) yields the same probability distribution over outcomes as the profile (f, g) .*

As we shall see later, (this version of) the revelation principle is useful but its usefulness is limited for a number of reasons. The first reason is that, although it restricts the class of mechanisms over which we must search to find the designer’s most preferred outcome, we still have to find the optimal device in this class, which is not always an easy task. The second reason is that in practice there will often be physical limitations on the devices that the designer can employ (because of limits to the device’s monitoring capabilities, for instance) and hence limitations on the communication mechanisms that should be considered, but these may not translate into limitations on a corresponding direct mechanism. For instance, in a flow control scenario, it will often be the case that the device can observe total flow but not the flow of individual users and can only observe this flow with errors; no such restrictions occur in direct mechanisms. Finally, as noted before, the revelation principle does not hold when communication is costly.

C. Special cases

The framework we have described is quite general so it is worth noting that many, perhaps more familiar, frameworks are simply special cases:

- If T , R , M and X are all singletons, then our framework reduces to an ordinary game in normal form and our equilibrium notion reduces to Nash equilibrium.
- If R , M and X are all singletons, then our framework reduces to an ordinary Bayesian game and our equilibrium notion reduces to Bayesian Nash equilibrium.
- If T , R and X are all singletons, then our framework reduces to a game with a mediation device and our equilibrium notion reduces to correlated equilibrium [33].
- If T , R and M are all singletons, then our framework reduces to the intervention framework of [3] and our equilibrium notion reduces to intervention equilibrium.
- If there are no errors, and reports and messages are costless, then our framework reduces to a communication game in the sense of [12] and our equilibrium notion reduces to communication equilibrium.

III. WHY INTERVENTION AND INFORMATION REVELATION MATTER

To illustrate our framework, we give a simple example to show that strategic behavior matters, intervention matters, and communication plus intervention matters – in the sense that they all change the outcomes that can be achieved.

We consider the problem of access to two channels A and B (e.g., two different bandwidths, or two different time slots). In each session, two users (identified as user 1 and user 2, but drawn from the same pool of users) can access either or both channels; we use A, B, AB to represent the obvious actions. Each user seeks to maximize its utility, which is the sum of its own goodput in the two channels.

Potential users are of four types: HL, ML, LM and LH ; the probability that a user is of a given type is $1/4$. We interpret a user's type xy as the quality of channels A, B to that user: channel A has quality x (Low, Medium or High), channel B has quality y (Low, Medium or High).¹² The goodput obtained by user $i = 1, 2$ from a given channel depends on the user's type and on which user(s) access the channel.

- if user i does not access the channel it obtains goodput = 0
- if both users access the channel they interfere with each other and both users obtain goodput = u_I
- if user i is the only user to access channel A and its type is xy then it obtains goodput u_x (where $x = L, M, H$)
- if user i is the only user to access channel B and its type is xy then it obtains goodput u_y (where $y = L, M, H$)

We assume $2u_I < u_L < u_M < u_H$.

We consider five scenarios: (I) no intervention or communication, (II) communication but no intervention, (III) intervention but no communication, (IV) intervention and communication, and (V) the benchmark setting in which the designer has perfect information and users are obedient. For simplicity, we assume that the devices available to the designer are very restricted: reports and messages are costless, there are no errors and the actions are either $x^* =$ "take no action" or $x_1 =$ "access both channels". If the device takes no action, user utilities are as above; if the device accesses both channels then each user's goodput is u_I on each channel the user accesses.¹³ The designer is benevolent and hence seeks to maximize social utility – the expected sum of user utilities.

I No communication, No Intervention Independently of the user's type, the other user's type, and the other user's action, it is always strictly better for each user to access both channels, so in the unique (Bayesian Nash) equilibrium both users always choose action AB , and (in obvious and suggestive notation), (expected) social utility is $EU(I) = 4u_I$.

II Communication, No Intervention Nothing changes from scenario I : no matter what the users report and the device recommends, it is strictly better for each user to access both channels, so in the unique equilibrium both users always choose action AB , and social utility is $EU(II) = 4u_I$.

¹²Note that the quality to a user is correlated across channels: each user finds one channel to be of Low quality and the other to be of Medium or High quality. This is not at all essential – the qualitative comparisons would be unchanged if we assumed quality to a user was uncorrelated across channels – but the calculation would be much messier.

¹³Note that in this model the utility obtained when accessing a channel in the presence of interference does not depend on the number of interferers present and on their channel qualities, which may not be realistic in certain scenarios. This assumption is made here in order to keep the discussion simple, but could be easily relaxed at the price of a much more cumbersome discussion in terms of notation and number of cases to be considered. In addition, in most reasonable scenarios (i.e., when the goodput obtained in the presence of any amount of interference is significantly lower than that obtained in its absence), the qualitative conclusions we draw here would be maintained.

III Intervention, No Communication The sets R_i of reports and M of messages are singletons, so the device obtains no information about the users and can suggest no actions to the users. The best the designer can do is to use an intervention rule that coordinates the two users to different resources; given the restriction on device actions an optimal rule is:

$$\Phi(a_1, a_2) = \begin{cases} x^* & \text{if } a_1 = A \text{ and } a_2 = B \\ x_1 & \text{otherwise} \end{cases}$$

where a_1 and a_2 are the actions adopted by the two users, which are assumed to be aware of the intervention rule. Given this intervention rule, the best equilibrium strategy profile (i.e., the one that yields highest social utility) is for user 1 to access channel A and user 2 to access channel B , so that there is never a conflict.¹⁴ Given the distribution of types, social utility is $EU(III) = (u_H/2) + (u_M/2) + u_L$.

IV Communication, Intervention We consider a direct mechanism in which the users report their types ($R_i = T_i$) and the device D recommends actions ($M = A$). The device uses the following message and intervention rules:

$$\mu(r_1, r_2) = \begin{cases} (A, B) & \text{if } r_1 = HL \text{ or } r_1 = ML \\ & \text{or } r_2 = LH \text{ or } r_2 = LM \\ (B, A) & \text{otherwise} \end{cases}$$

$$\Phi(r_1, r_2, a_1, a_2) = \begin{cases} x^* & \text{if } (a_1, a_2) = \mu(r_1, r_2) \\ x_1 & \text{otherwise} \end{cases}$$

where r_1, r_2 are the reports and a_1, a_2 are the actions. This is an incentive compatible direct mechanism. To see this we must show that the honest and obedient strategy (f_1^*, g_1^*) is the most preferred strategy for all types of user 1, given that user 2 follows its honest and obedient strategy (f_2^*, g_2^*) , and conversely for user 2. We will describe the calculations for user 1, from which those for user 2 can be derived by the symmetry of the problem.

Assume user 1 is of type HL . If it is honest and obedient, it obtains a utility of u_H because it accesses its preferred channel. This utility is always higher than the utility it obtains not being obedient, i.e., if it does not follow the recommendation. In fact in this case it never obtains a utility higher than $2u_I$ because the channels are interfered by the device. Now assume user 1 is obedient but not honest. If it reports type ML it can still access its preferred channel, obtaining a utility of u_H , the same as if it were honest. If it reports type LM or LH , it accesses half of the time its preferred channel and half of the time its less preferred channel (depending on the type of user 2), obtaining an expected utility of $(u_H + u_L)/2$ which is lower than u_H . These considerations translate mathematically into the set of inequalities in Fig. 1 for $EU_1(f^*, g^*, HL, D)$, stating that user 1 has an incentive to be honest and obedient if it is of type HL . Fig. 1 also shows the analogous inequalities if user 1 is of type ML, LM or LH .

Notice that following this mechanism never leads to interference (users always access different channels) and users are "assigned" to the most efficient channels $7/8$ of the time. However, users are not *always* assigned to the most efficient channels: if type profiles are $(t_1, t_2) = (ML, HL)$

¹⁴This is not the only equilibrium but it is the best, both for the designer and for the users. In the other equilibrium the users access both channels.

$$\begin{aligned}
EU_1(f^*, g^*, HL, D) &= \begin{cases} u_H > 2u_I \geq EU_1(f_1, f_2^*, g_1, g_2^*, HL, D) & \forall f_1, \text{ if } g_1(HL, a_1) \neq a_1 \\ u_H = EU_1(f_1, f_2^*, g^*, HL, D) & \text{if } f_1(HL) = ML \\ u_H > (u_H + u_L)/2 = EU_1(f_1, f_2^*, g^*, HL, D) & \text{if } f_1(HL) = LM \text{ or } LH \end{cases} \\
EU_1(f^*, g^*, ML, D) &= \begin{cases} u_M > 2u_I \geq EU_1(f_1, f_2^*, g_1, g_2^*, ML, D) & \forall f_1, \text{ if } g_1(ML, a_1) \neq a_1 \\ u_M = EU_1(f_1, f_2^*, g^*, ML, D) & \text{if } f_1(ML) = HL \\ u_M > (u_H + u_L)/2 = EU_1(f_1, f_2^*, g^*, ML, D) & \text{if } f_1(ML) = LM \text{ or } LH \end{cases} \\
EU_1(f^*, g^*, LM, D) &= \begin{cases} (u_M + u_L)/2 > 2u_I \geq EU_1(f_1, f_2^*, g_1, g_2^*, LM, D) & \forall f_1, \text{ if } g_1(LM, a_1) \neq a_1 \\ (u_M + u_L)/2 = EU_1(f_1, f_2^*, g^*, LM, D) & \text{if } f_1(LM) = LH \\ (u_M + u_L)/2 > u_L = EU_1(f_1, f_2^*, g^*, LM, D) & \text{if } f_1(LM) = HL \text{ or } ML \end{cases} \\
EU_1(f^*, g^*, LH, D) &= \begin{cases} (u_H + u_L)/2 > 2u_I \geq EU_1(f_1, f_2^*, g_1, g_2^*, LH, D) & \forall f_1, \text{ if } g_1(LH, a_1) \neq a_1 \\ (u_H + u_L)/2 = EU_1(f_1, f_2^*, g^*, LH, D) & \text{if } f_1(LH) = LM \\ (u_H + u_L)/2 > u_L = EU_1(f_1, f_2^*, g^*, LH, D) & \text{if } f_1(LH) = HL \text{ or } ML \end{cases}
\end{aligned}$$

Fig. 1. Set of inequalities showing that user 1 has an incentive to be honest and obedient.

or $(t_1, t_2) = (LH, LM)$ then user 1 is assigned to channel A and user 2 is assigned to channel B , which is inefficient. This inefficiency is an unavoidable consequence of incentive compatibility: if user 2 were always assigned to his preferred channel A when he reported HL (for instance) then he would never be willing to report ML when that was his true type. Expected social utility under this mechanism is

$$EU(IV) = (3u_H/4) + (3u_M/4) + (u_L/2)$$

V Benchmark Social Optimum: Public Information, Perfect Cooperation The social optimum is obtained by assigning the user with the best channel quality to his favorite channel and never assigning two users to the same channel. Expected social utility is

$$EU(V) = (7u_H/8) + (5u_M/8) + (u_L/2)$$

Direct calculation shows that social utilities in four of the five scenarios are strictly ranked:

$$EU(I) = EU(II) < EU(III) < EU(IV) < EU(V)$$

In words: in comparison to the purely Bayesian scenario (no intervention), communication without intervention achieves nothing, intervention without communication improves social utility by dampening destructive competition, intervention with communication improves social utility even more by extracting some information and using that information to promote a more efficient coordination across types, but even intervention with communication does not achieve the benchmark social optimum under full cooperation. It is possible to show that the same conclusions would be obtained in an environment with n users and m channels (for arbitrary n, m), provided that $m \geq n$ and $mu_I < u_L < u_M < u_H$.

It is worth noting that similar comparisons across scenarios could be made in many environments and the ordering of expected social utility would be as above:

$$EU(I) \leq EU(II) \leq EU(III) \leq EU(IV) \leq EU(V)$$

In general, any of these inequalities might be strict.

IV. RESOURCE ALLOCATION GAMES IN COMMUNICATION ENGINEERING

In the following sections of the paper, we explore the designer's problem in a class of abstract environments that

exhibit some features common to many resource sharing situations in communication networks, including power control [6], [13], MAC, [4], [14], and flow control [14]–[18].

The contributions of this part of the paper are the following. We characterize the direct communication mechanisms that are optimal among all mechanisms. We provide conditions on the environment under which it is *possible* for the designer to achieve its benchmark optimum – the outcome it could achieve if users were compliant – and conditions under which it is *impossible* for the designer to achieve its benchmark optimum. Although we can characterize the optimal device, other mechanisms are also of interest, for several reasons. The optimal device may be very difficult to compute. It is therefore of some interest to consider mechanisms that are sub-optimal but easy to compute, and we provide a simple algorithm that converges to such a mechanism. Moreover, in some situations, it may not be possible for the users to communicate with the designer, so it is natural to consider intervention schemes that do not require the users to make reports. These sub-optimal devices are very useful to implement practical schemes and, for this reason, are used in [19] to design a flow control management system robust to self-interested and strategic users.

A. The considered environment

In this subsection we formalize the particular (but, at the same time, quite general) environment we consider from now on, motivating each assumption with examples of its application in resource sharing situations in communication networks.

We consider a finite and discrete type set made by real numbers $T_i = \{\tau_{i,1}, \tau_{i,2}, \dots, \tau_{i,v_i}\} \subset \mathbb{R}$, $v_i \in \mathbb{N}$, in which the elements are labeled in increasing order, $\tau_{i,1} < \tau_{i,2} < \dots < \tau_{i,v_i}$. We interpret the type of a user as the valuation of a particular resource for the user (e.g., different types may represent different quality of service classes). We assume that every type profile has a positive probability to occur, i.e., $\pi(t) > 0$, $\forall t$. We allow the users to take actions in a continuous interval $A_i = [a_i^{min}, a_i^{max}] \subset \mathbb{R}$, which we interpret as the level of resource usage (e.g., it may represent the adopted transmission power, which is positive and upper bounded). We assume that the devices available to the designer are such

that reports and messages are costless, there are no errors, and there exists the intervention action $x^* \in X$ which we interpret as “no intervention”. In this case we can simply write $U_i(a, t, x)$ for the utility of user i and $U(a, t, x)$ for the utility of the designer and we can restrict our attention to incentive compatible direct mechanisms. That is, we consider only the incentive compatible devices $D = \langle (T_i), (A_i), \mu, X, \Phi \rangle$ in which $x^* \in X$.

We assume that the designer’s utility satisfies the following assumptions, $\forall t \in T$,

A1: $U(a, t, x^*) > U(a, t, x)$, $\forall a \in A$, $\forall X$, $\forall x \in X$, $x \neq x^*$

A2: $g^M(t) = \operatorname{argmax}_a U(a, t, x^*)$ is unique

A3: $g^M(t)$ is differentiable with respect to t_i and $\frac{\partial g_i^M(t)}{\partial t_i} > 0$ ¹⁵

Assumption **A1** states that the “no intervention” action is the strictly preferred action of the designer, regardless of users’ actions and types. Interpreting interventions as punishments, assumption **A1** asserts that the designer is not happy if the users are punished.

Assumption **A2** states that, for every type profile $t \in T$, the users’ joint action profile that maximizes the designer’s utility is unique, and by assumption **A3**, each component in g^M is continuous and increasing in the type of that user. If actions represent the level of resource usage and types represent resource valuations, assumption **A3** asserts that the higher i ’s valuation the higher i ’s socially optimum level of resource usage.

Under these assumptions, the benchmark optimum for the designer can be easily determined

$$EU^{ben} = \sum_{t \in T} \pi(t) U(g^M(t), t, x^*) \quad (2)$$

For each type profile $t \in T$, (we define the complete information game

$$\Gamma_t^0 = (N, A, \{U_i(\cdot, t, x^*)\}_{i=1}^n)$$

Γ_t^0 is the complete information game (users know everybody’s type) that can be derived from our general framework assuming that sets of types T_i , reports R_i , messages M_i and interventions X are singletons (in particular, $X = \{x^*\}$). It can be thought as the game that models users’ interaction in the absence of an intervention device and when the type profile is known.

The strategy of user i in this context is represented by the function $g_i : T \rightarrow A_i$ (notice that we can omit the dependence on the messages), since the function $f : T \rightarrow R$ is automatically defined (users do not send reports or receive messages, or equivalently, always send the report “no report” and receive the message “no message”). We denote by $g^{NE^0}(t) = (g_1^{NE^0}(t), \dots, g_n^{NE^0}(t))$ a Nash Equilibrium (NE) of the game Γ_t^0 , which is an action profile so that each user obtains its maximum utility given the actions of the other users, i.e., $\forall i \in N$ and $\forall g_i : T \times \{m^*\} \rightarrow A_i$

$$U_i(g^{NE^0}(t), t, x^*) \geq U_i(g_i(t), g_{-i}^{NE^0}(t), t, x^*)$$

¹⁵This assumption requires the designer’s utility to be defined over a continuous interval that includes the finite type set T . Note that if $\frac{\partial g_i^M(t)}{\partial t_i} < 0$ it is possible to define a new type variable as the inverse of the original type, such that **A3** is satisfied.

We assume that users’ utilities $U_i(a, t, x^*)$ are twice differentiable¹⁶ with respect to a and, $\forall a \in A$, $\forall t \in T$, $\forall i, j \in N$, $i \neq j$,

A4: $U_i(a, t, x^*)$ is quasi-concave in a_i and there exists a unique best response function $h_i^{BR}(a_{-i}, t) = \operatorname{argmax}_{a_i} U_i(a, t, x^*)$

A5: $\frac{\partial^2 U_i(a, t, x^*)}{\partial a_i \partial a_j} \leq 0$

A6: There exists g^{NE^0} such that $g^{NE^0}(t) \geq g^M(t)$ ¹⁷ and $g_k^{NE^0}(t_k, t_{-k}) > g_k^M(t_k, t_{-k})$ for some user $k \in N$ and type $t_k \in T_k$

Since for **A4** the users’ utilities are quasi concave (thus the game Γ_t^0 is a quasi-concave game) and the best response function $h_i^{BR}(a_{-i}, t)$ that maximizes $U_i(a, t, x^*)$ is unique, either i ’s utility is monotonic with respect to a_i , or it increases with a_i until it reaches a maximum for $h_i^{BR}(a_{-i}, t)$, and decreases for higher values. As a consequence, a NE $g^{NE^0}(t)$ of Γ_t^0 exists. In fact, the best response function $h^{BR}(a, t) = (h_1^{BR}(a_{-1}, t), \dots, h_n^{BR}(a_{-n}, t))$ is a continuous function from the convex and compact set A to A itself, therefore Brouwer’s fixed point theorem ensures that a fixed point exists.

Assumption **A5** asserts that Γ_t^0 is a submodular game and ensures that $h_i^{BR}(a_{-i}, t)$ is a non increasing function of a_j , $j \neq i$. Interpreting a_i as i ’s level of resource usage, this situation reflects resource allocation games where it is in the interest of a user not to increase its resource usage if the total level of use of the other users increases, in order to avoid an excessive use of the resource. Nevertheless, assumption **A6** says that strategic users use the resources more heavily compared to the optimal (from the designer’s point of view) usage level.

The class of games satisfying assumptions **A1-A6** includes the linearly coupled games [14] and many resource allocation games in communication networks, such as the MAC [4], [14], power control [6], [13] and flow control [14]–[18] games, assuming that the designer’s utility is increasing in the users’ utilities (i.e., a benevolent designer).

B. Intervention in the complete information setting

As a further step toward the objective of this paper which is the design of a mechanism dealing with both information revelation and action enforcement, we introduce here the special case of intervention in the complete information setting.

For each type profile $t \in T$ and intervention rule $\Phi : A \rightarrow \Delta(X)$ (we can omit the dependence on reports and messages), we define the complete information game

$$\Gamma_t = (N, A, \{U_i(\cdot, t, x)\}_{i=1}^n)$$

where x is drawn according to the distribution $\Phi(a)$, i.e., with probability $\Phi_a(x)$. Γ_t is the complete information game (users and designer know everybody’s type) that can be derived from our general framework assuming that sets of types T_i , reports R_i and messages M_i are singletons. Our general framework

¹⁶All the results go through if utility functions are only continuous. The assumption that utility functions are twice differentiable is made only in order that submodularity (**A5**) takes a particularly convenient form.

¹⁷Throughout the paper, inequalities between vectors are intended component-wise.

reduces in this case to the intervention framework of [3] and our equilibrium notion reduces to intervention equilibrium.

As in the game Γ_t^0 , the strategy of user i is represented only by the function $g_i : T \rightarrow A_i$, but in this case i has to take into account the effect on its (expected) utility of the intervention action x , which depends on the adopted action profile a . According to the notions introduced in the general framework, we say that a device D , defined by the set of interventions X and the intervention rule Φ , sustains (without intervention) the strategy profile $g(t) = (g_1(t), \dots, g_n(t))$ in Γ_t if g is an equilibrium of Γ_t (and $\Phi_{g(t)}(x^*) = 1$). If there exists a device D able to sustain (without intervention) the profile g in Γ_t , we say that g is sustainable (without intervention) in Γ_t .

V. OPTIMAL DEVICES

In this section we study the class of environments introduced in Section IV with the general framework proposed in Section II. In particular, we take the part of a designer seeking to maximize his own expected utility in the presence of self-interested and strategic users, choosing an optimal device in the class of available devices \mathcal{D} specified in Section IV.

First of all we wonder if the designer can choose a maximum efficiency device $D \in \mathcal{D}$ to obtain his benchmark utility despite the fact that the users are strategic. We characterize the existence and the computation of maximum efficiency devices based on some properties of the complete information setting. Moreover, we prove that a necessary condition for the existence of a maximum efficiency device requires the type sets to be *sufficiently sparse*.

Even for cases in which a maximum efficiency device does not exist, the designer is still interested in obtaining the best he can, choosing an optimal device. For this reason we study the problem of finding the optimal device and we prove that, under some properties of the complete information setting, the original problem can be decoupled into two sub-problems that are easier to solve.

A. Properties of a maximum efficiency device

In this subsection we address the problem of the existence and the computation of a maximum efficiency incentive compatible device.

The first result we derive asserts that a maximum efficiency device exists if and only if, for every type profile t , the optimal (for the designer) strategy profile $g^M(t)$ is sustainable without intervention in Γ_t , and users have incentives to reveal their real type given that they will adopt g^M and the intervention device does not intervene. If this is the case, we are also able to characterize all maximum efficiency devices.

Proposition 2. $D = \langle (T_i), (A_i), \mu, X, \Phi \rangle$ is a maximum efficiency device if and only if, $\forall t \in T$,

- 1: the optimal action profile $g^M(t)$ is sustainable without intervention in Γ_t ;
- 2: each user i having type t_i prefers the action profile $g^M(t)$ with respect to the action profile $g^M(t'_i, t_{-i})$, for every

alternative type t'_i user i might have, i.e.,

$$\begin{aligned} \sum_{t_{-i} \in T_{-i}} \pi(t | t_i) U_i(g^M(t), t, x^*) &\geq \\ &\geq \sum_{t_{-i} \in T_{-i}} \pi(t | t_i) U_i(g^M(t'_i, t_{-i}), t, x^*) \\ \forall i \in N, \forall t_i \in T_i, \forall t'_i \in T_i \end{aligned}$$

- 3: the suggested action profile is the optimal action profile of game Γ_t , i.e., $\mu(t) = g^M(t)$;
- 4: the restriction of the intervention rule in $r = t$ and $m = g^M(t)$, i.e., $\Phi_a = \Phi_{t, g^M(t), a}$, sustains without intervention $g^M(t)$ in Γ_t .

Proof: See Appendix A ■

Condition 1 is related to what is achievable by the designer in the complete information setting, condition 2 is related to the structure of the environment (which is not controllable by the designer), while conditions 3-4 say how to obtain a maximum efficiency direct mechanism once 1-2 are satisfied.

In the second result we combine condition 2 of Proposition 2 with assumptions **A3-A6** to derive a sufficient condition on the type set structures under which a maximum efficiency incentive compatible direct mechanism does not exist. We define the bin size β_k of user k 's type set, T_k , as the maximum distance between two consecutive elements of T_k : $\beta_k = \max_{s \in \{1, \dots, v_k - 1\}} (\tau_{k, s+1} - \tau_{k, s})$. We define the bin size β as the maximum among the bin sizes of all users: $\beta = \max_{k \in N} \beta_k$.

Proposition 3. There exists a threshold bin size $\zeta > 0$ so that if $\beta \leq \zeta$ then a maximum efficiency incentive compatible direct mechanism does not exist.

Proof: Let $k \in N$ and $t_k \in T_k$ be such that $g_k^{NE^0}(t) > g_k^M(t)$, $\forall t_{-k} \in T_{-k}$. We rewrite condition 2 of Proposition 2 for user k and type t_k :

$$\begin{aligned} \sum_{t_{-k} \in T_{-k}} \pi(t | t_k) U_k(g^M(t_k, t_{-k}), t, x^*) &\geq \\ &\geq \sum_{t_{-k} \in T_{-k}} \pi(t | t_k) U_k(g^M(t'_k, t_{-k}), t, x^*) \end{aligned} \quad (3)$$

$\forall t'_k \in T_k$. We have $h_k^{BR}(g_{-k}^M(t), t) \geq h_k^{BR}(g_{-k}^{NE^0}(t), t) = g_k^{NE^0}(t) > g_k^M(t)$, where the first inequality is valid because of the submodularity.

Let $\tilde{t}_k(t_{-k})$ be the value of user k 's type so that $g_k^M(\tilde{t}_k(t_{-k}), t_{-k}) = h_k^{BR}(g_{-k}^M(\tilde{t}_k(t_{-k}), t_{-k}), t)$ if it exists (in this case **A3** guarantees it is greater than t_k); and $\tilde{t}_k(t_{-k}) = \tau_{k, v_k}$ otherwise. Let $\hat{t}_k = \min_{t_{-k}} \tilde{t}_k(t_{-k})$. If $(t_k, \hat{t}_k] \cap T_k \neq \emptyset$ (in particular, this is true if $\beta \leq \hat{t}_k - t_k$), $\forall t'_k \in (t_k, \hat{t}_k] \cap T_k$ and $\forall t_{-k} \in T_{-k}$ we obtain

$$U_k(g^M(t'_k, t_{-k}), t, x^*) > U_k(g^M(t_k, t_{-k}), t, x^*)$$

contradicting Eq. (3). ■

Interpretation: when user k 's type is t_k , k 's resource usage that maximizes the designer's utility, $g_k^M(t)$, is lower than the one that maximizes k 's utility, $h_k^{BR}(g_{-k}^M(t), t)$, $\forall t_{-k} \in T_{-k}$. If k reports a type t'_k slightly higher than t_k , then the intervention device suggests a slightly higher resource usage,

allowing k to obtain a higher utility. Hence, k has an incentive to cheat and resources are not allocated as efficiently as possible. To avoid this situation, the intervention device might decrease the resources given to a type t'_k . In this case the loss of efficiency occurs when the real type of k is t'_k and it does not receive the resources it would deserve. These two cases are such that at least one of them corresponds to a non-zero inefficiency. Since both occur with positive probability, a positive overall inefficiency is unavoidable.

It is worth noting that we consider finite type sets and a finite intervention rule set mainly to simplify the logical exposition. However, all results might be derived also with infinite and continuous sets. In particular, if type sets are continuous Proposition 3 implies that a maximum efficiency incentive compatible direct mechanism never exists.

B. Properties of an optimal device

If a maximum efficiency device exists, the set of optimal devices in \mathcal{D} coincides with the set of maximum efficiency devices in \mathcal{D} , that is characterized in Proposition 2. If a maximum efficiency device does not exist, the designer seeks to obtain the best he can, minimizing the loss of efficiency. He has to choose the optimal device solving the **OD** problem. However, this may be computationally hard.

In this subsection we consider some additional conditions to simplify the **OD** problem. First, we assume that the designer's utility is a function of the users' utilities (this is the case, for example, of a benevolent designer that seeks to maximize some measure of social welfare). Moreover, we suppose that, for each type profile $t \in T$, every action profile $g(t)$ lower than $g^{NE^0}(t)$ is sustainable without intervention in Γ_t . Finally, we assume that the utility of a user i adopting the lowest action a_i^{min} is always equal to 0, i.e., $U_i(a_i^{min}, a_{-i}, t, x) = 0$, $\forall a_{-i}, t, x$. Interpreting a_i^{min} as no resource usage, this means that, independently of types and other users' actions, a user that does not use resources obtains no utility. In particular, this last assumption implies that:

Lemma 1. *The utility of user i is non increasing in the actions of the other users.*

Proof:

$$\begin{aligned} U_i(a, t, x) &= U_i(a_i^{min}, a_{-i}, t, x) + \int_{a_i^{min}}^{a_i} \frac{\partial U_i(z, a_{-i}, t, x)}{\partial z} \partial z \\ &= \int_{a_i^{min}}^{a_i} \frac{\partial U_i(z, a, t, x)}{\partial z} \partial z \\ \frac{\partial U_i(a, t, x)}{\partial a_j} &= \int_{a_i^{min}}^{a_i} \frac{\partial^2 U_i(z, a_{-i}, t, x)}{\partial z \partial a_j} \partial z \leq 0 \end{aligned}$$

where the inequality is valid because of **A5**. \blacksquare

Under the additional assumptions of this subsection, we can prove the following result that allows the designer to further restrict the class of mechanisms to take into consideration.

Lemma 2. *There exists an optimal device such that, for every type profile $t \in T$, the recommended action profile $\tilde{a}(t)$ is unique (i.e., μ is point mass at $\tilde{a}(t)$) and the restriction of the intervention rule in $r = t$ and $m = \tilde{a}(t)$, i.e., $\Phi_a = \Phi_{t, \tilde{a}(t), a}$, sustains without intervention $\tilde{a}(t)$ in Γ_t .*

Proof: See Appendix B \blacksquare

Lemma 2 suggests the idea to decouple the original problem into two sub-problems. First, we can calculate the optimal message rule $\tilde{\mu}$ under the constraint that users adopting the recommended actions have the incentive to report their real type. Then, it is sufficient to identify an intervention rule $\tilde{\Phi}$ able to sustain $\tilde{\mu}(t)$ without intervention in Γ_t , $\forall t$. This is formalized in the following.

Consider the device $\tilde{D} = \langle (T_i), (A_i), \tilde{\mu}, X, \tilde{\Phi} \rangle$, where

$$\tilde{\mu} = \operatorname{argmax}_{\mu} \sum_{t \in T} \pi(t) U(\mu(t), t, x^*)$$

subject to:

$$\sum_{t_{-i} \in T_{-i}} \pi(t | t_i) U_i(\mu(t), t, x^*) \geq \sum_{t_{-i} \in T_{-i}} \pi(t | t_i) U_i(\mu(t'_i, t_{-i}), t, x^*)$$

$$\forall i \in N, \forall t_i \in T_i, \forall t'_i \in T_i$$

and, $\forall t \in T$, $\Phi_a = \tilde{\Phi}_{t, \mu(t), a}$ sustains without intervention $\mu(t)$ in Γ_t .

Proposition 4. *\tilde{D} is an optimal device.*

Proof: Lemma 2 guarantees that there exists an optimal device inside the class of devices in which, $\forall t$, the recommended action profile $\mu(t)$ is unique and the restriction of the intervention rule in $r = t$ and $m = \mu(t)$, i.e., $\Phi_a = \Phi_{t, \mu(t), a}$, sustains without intervention $\mu(t)$ in Γ_t . Among all devices belonging to such class, \tilde{D} is selected to maximize the designer's expected utility. Thus, \tilde{D} is an optimal device. \blacksquare

VI. ALGORITHM THAT CONVERGES TO AN INCENTIVE COMPATIBLE DEVICE

In this section we provide a practical tool for the designer to choose an efficient device. Because the optimal device may be very difficult to compute, even in the decoupled version of Proposition 4, we provide a simple algorithm that converges to an incentive compatible device D in which μ is point mass (i.e., given a report the recommended action profile is unique) and, although perhaps not optimal, still yields a "good" outcome for the designer. More precisely, D will sustain *without intervention* the honest and obedient strategy profile. The algorithm has been designed with the idea to minimize the distance between the optimal action profile $g^M(t)$ and the suggested action profile $\mu(t)$, for each possible type profile t . Such an algorithm is run by the designer to choose an efficient device and can be used when, for every type profile t and at each step of the algorithm, the designer is able to identify a device for the complete information setting that sustains without intervention the suggested action profile $\mu(r)$ in Γ_r . (Note that the suggested action profile will never be lower than the optimal action profile $g^M(t)$ or higher than the NE action profile $g^{NE^0}(t)$ of Γ_i^0 .)

Given a device D in which μ is point mass, we denote by $W_i(t_i, t'_i)$ the expected utility that user i , with type t_i , obtains reporting type t'_i and adopting the suggested action, when the other users are honest and obedient and the intervention device does not intervene, i.e.,

$$W_i(t_i, t'_i) = \sum_{t_{-i} \in T_{-i}} \pi(t | t_i) U_i(\mu(t'_i, t_{-i}), t, x^*)$$

Moreover, we say that X and $\Phi_{r,m,a}$ are *induced* by μ if the device defined by X and $\Phi_a(x) = \Phi_{r,\mu(r),a}(x)$ sustains $\mu(r)$ without intervention in $\Gamma_r, \forall r \in T$. If the designer is able to identify the device defined by X and Γ_r in the complete information setting, then he can easily compute X and $\Phi_{r,m,a}$ induced by μ , obtaining a device for the general framework that, by construction, gives the users the incentive to always adopt the recommended actions (i.e., users are obedient) and does not intervene (threats of punishments do not need to be executed since users follow the recommendations).

The algorithm initializes the device D in the following way: $\mu(r) = g^M(r)$, X and Φ induced by μ . This means that, given the report profile r , the device recommends the optimal (for the designer and if user types are r) action profile $g^M(r)$ and the users will adopt it. However, this does not guarantee that the users are honest: the reported type profile may be different from the real one, i.e., $r \neq t$. To give an incentive for the users to be honest, in each step of the algorithm the recommended action profile $\mu(r)$ is modified to increase the utility the users obtain if they are honest (or to decrease the utility they obtain when they are dishonest). Whenever $\mu(r)$ is modified, also X and Φ must be modified accordingly, selecting X and Φ induced by μ such that users remain obedient.

To explain the idea behind the algorithm we exploit Fig. 2, where i 's utility is plotted with respect to i 's action, for a fixed type profile t , when all users are honest (i.e., $r = t$) and the other users are obedient (i.e., $g_j(t_j, \mu_j(t)) = \mu_j(t), \forall j \neq i$). Each sub-picture refers to different recommended actions (i.e., different μ), and in each sub-picture four points are marked (some of which may possibly coincide) representing the following cases: (1) i adopts the best (for the designer) action $g_i^M(t)$; (2) i adopts the recommended action $\mu_i(t)$; (3) i adopts the NE action $g_i^{NE^0}(t)$ (notice that it is not the best action for user i because the other users do not adopt $g_i^{NE^0}(t)$); and (4) i adopts the best action $h_i^{BR}(\mu_{-i}(t), t)$.

The initialization case, in which (1) and (2) coincide, is represented by the upper-left Fig. 2. By assumption **A6** $g_i^M(t) \leq g_i^{NE^0}(t)$ and by assumption **A5** $g_i^{NE^0}(t) \leq h_i^{BR}(\mu_{-i}(t), t)$, because $\mu_{-i}(t) \leq g_{-i}^{NE^0}(t)$. If $W_i(t_i, t_i) \geq W_i(t_i, t'_i)$, for every alternative i 's reported type t'_i , then user i has an incentive to report its true type t_i . If, at a certain iteration of the algorithm, this is valid for all users and for all types they may have, then the algorithm stops and we obtain a device that sustains without intervention the honest and obedient strategy profile.¹⁸

Conversely, suppose there exists a user i and types t_i and t'_i such that $W_i(t_i, t_i) < W_i(t_i, t'_i)$, i.e., user i has an incentive to report t'_i when its type is t_i . In this case the algorithm increases the recommended action $\mu_i(t)$ by a quantity equal to ϵ_i , moving it in the direction of the best response function $h_i^{BR}(\mu_{-i}(t), t)$, for every possible combination of types t_{-i} of the other users, and X and Φ must be modified accordingly such that users remain obedient. This has the effect, as represented by upper-right Fig. 2, to increase the utility of user i when it is honest, $\forall t_{-i}$, which in turn implies that the expected utility of users i when it is honest (i.e., $W(t_i, t_i)$) increases.

¹⁸Notice that, if a maximum efficiency incentive compatible direct mechanism exists, since it must satisfy the conditions of Proposition 2, then the initialization of the algorithm corresponds to a maximum efficiency incentive compatible direct mechanism and the algorithm stops after the first iteration.

This procedure is repeated as long as $W_i(t_i, t_i) < W_i(t_i, t'_i)$ and $\mu_i(t) \leq g_i^{NE^0}(t)$.

In case i 's suggested action $\mu_i(t)$ reaches $g_i^{NE^0}(t)$ and still $W_i(t_i, t_i) < W_i(t_i, t'_i)$, then the suggested action of user k , $\mu_k(t)$, is increased by a quantity equal to $\epsilon_k, \forall k \in N, k \neq i, \forall t_{-i} \in T_{-i}$. As we can see from lower-left Fig. 2, this means to change the shape of the curve representing i 's utility with respect to i 's action. In particular, by assumption **A5**, the best response function $h_i^{BR}(\mu_{-i}(t), t)$ is moved in the direction of the recommended action $\mu_i(t)$.

If $\mu_k(t)$ reaches $g_k^{NE^0}(t)$ as well, $\forall k \in N$, then $\mu_i(t)$ coincides with the best response function $h_i^{BR}(\mu_{-i}(t), t)$, as represented in the lower-right Fig. 2. In fact, by definition, the NE is the action profile such that every user is playing its best action against the actions of the other users. Since $\mu_i(t)$ coincides with $h_i^{BR}(\mu_{-i}(t), t), \forall t_{-i} \in T_{-i}$, user i is told to play its best action for every possible combination of the types of the other users. Hence, user i cannot increase its utility reporting type t'_i , i.e., it must be $W_i(t_i, t_i) \geq W_i(t_i, t'_i)$.

The algorithm stops the first time every user has an incentive to declare its real type. Since at each iteration the suggested action profiles are increased by a fixed amount, the algorithm converges after a finite number of iterations. The higher the steps $\epsilon_i, i \in N$, the lower the convergence time of the algorithm. On the other hand, the lower the steps, the closer the suggested action profile to the optimal one.¹⁹

Algorithm 1 General algorithm.

- 1: **Initialization:** $\mu(t) = g^M(t) \forall t, X$ and Φ induced by μ
 - 2: **For** each user $i \in N$ and each pair of types $t_i, t'_i \in T_i$
 - 3: **If** $W_i(t_i, t_i) < W_i(t_i, t'_i)$
 - 4: **If** $\mu_i(t) < g_i^{NE^0}(t)$ for some $t_{-i} \in T_{-i}$
 - 5: $\mu_i(t) \leftarrow \min\{\mu_i(t) + \epsilon_i, g_i^{NE^0}(t)\}, \forall t_{-i} \in T_{-i}$
 - 6: X and Φ induced by μ
 - 7: **Else**
 - 8: $\mu_k(t) \leftarrow \min\{\mu_k(t) + \epsilon_k, g_k^{NE^0}(t)\}, \forall k \neq i, \forall t_{-i} \in T_{-i}$
 - 9: X and Φ induced by μ
 - 10: Repeat from 2 until 3 is unsatisfied $\forall i, t_i, t_{-i}$
-




Fig. 2. User i 's utility vs. user i 's action, for different recommended actions.

A. Example: a 2-users MAC game

In this subsection we provide a simple example to illustrate the concepts introduced in the last sections and to validate the

¹⁹Notice that, since no assumption such as convexity is made for the designer's expected utility, an action profile closer to the optimal one does not necessarily imply a better outcome for the designer.

usefulness of the proposed algorithm.

Consider a MAC problem in which two users, 1 and 2, contend for a common channel. Time is slotted, and in each slot user i , $i = 1, 2$, transmits a packet with probability $a_i \in [0, 1]$, which is chosen by user i at the beginning of the communication sessions and kept constant. Moreover, in the system there is an intervention device D that transmits packets with probability $x \in [0, 1]$. We assume that a packet is correctly received if and only if a collision does not occur. Hence, the rate of user i is $R_i = a_i(1 - a_{-i})(1 - x)$, where the subscript $-i$ denotes the other user. We consider the following utility functions for the users and the designer

$$U_i(a, t_i, x) = \log R_i - \frac{a_i}{t_i} + c_i, \quad i = 1, 2$$

$$U(a, t, x) = U_1(a, t_1, x) + U_2(a, t_2, x)$$

where $a = (a_1, a_2)$, $t = (t_1, t_2)$, $t_i \in \mathfrak{R}$ is the type of user i , and c_i is a constant.

The utility of user i is increasing and concave with respect to i 's throughput, and linearly decreasing with respect to i 's transmission probability, that represents i 's average power consumption per slot. The type t_i , which can be interpreted as the number of transmissions per unit of cost, expresses the trade-off between the relative importance of throughput and power for user i .

$x^* = 0$ represents the ‘‘no intervention’’ action. $U_i(a, t_i, 0)$ is concave with respect to a_i , and imposing $\frac{\partial U_i(a, t_i, 0)}{\partial a_i} = 0$, the resulting NE of the game Γ_0^t is $g_i^{NE^0}(t) = t_i$, $i = 1, 2$. Also, $U(a, t, 0)$ is jointly concave with respect to a_1 and a_2 , and imposing $\frac{\partial U(a, t, 0)}{\partial a_i} = 0$, the resulting optimal action is $g_i^M(t) = t_i + \frac{1}{2} - \sqrt{t_i^2 + \frac{1}{4}}$, $i = 1, 2$.

It is possible to check that in this case all the conditions **A1-A6** are satisfied. Moreover, also the additional assumptions made in Section V.B are satisfied. In particular, if the device's punishments are strong enough (i.e., x is high enough) when some users do not follow the recommendations, it is possible to show that every action profile $g(t) \leq g_i^{NE^0}(t)$ is sustainable without intervention in Γ_t (i.e., the complete information setting).

Now we consider some numerical examples to describe the results derived in Sections V and VI. We assume a common constant $c_1 = c_2 = 3$ (the analysis is independent of the constant, this choice is made only to obtain positive utilities), a common type set $T_1 = T_2 = \{t_{low}, t_{high}\}$, and a user is of a particular type with probability 0.5, independently of the other user's type. If $T_1 = T_2 = \{1/4, 5/6\}$, it is easy to check that condition 2 of Proposition 2 is satisfied, and in fact it is possible to adopt a maximum efficiency device which asks the users to report their types (they will be honest to avoid unfavorable recommendations) and recommends them to adopt the optimal actions (they will be obedient to avoid the punishments).

If $T_1 = T_2 = \{1/3, 5/6\}$, then a maximum efficiency mechanism does not exist. In fact, if the device recommended the users to adopt the optimal actions, an honest low type user would obtain an average utility of 0.84, whereas a dishonest low type user would obtain an average utility of

0.87, i.e., low type users have an incentive to cheat by declaring themselves to be of high type. To avoid this situation the proposed algorithm increases the transmission probabilities recommended to low type users. If $\epsilon = 0.01$, after 8 iterations (i.e., the recommended action for a low type user is increased by 0.08) the algorithm converges to an incentive compatible device in which the designer utility is 1.63, slightly lower than the benchmark optimum 1.68, but higher than 0.52, which is the utility obtainable in the absence of an intervention device.

VII. CONCLUSION

In this paper we have extended the intervention framework introduced by [3] to take into account situations of private information, imperfect monitoring and costly communication. We allow the designer to use a device that can communicate with users and intervene in the system. The goal of the designer is to choose the device that allows him to obtain the highest possible utility in the considered scenario. For a class of environments that includes many engineering scenarios of interest (e.g., power control [6], [13], MAC [4], [14], and flow control [14]–[18]) we find conditions under which there exist devices that achieve the benchmark optimum and conditions under which such devices do not exist. In case they do not exist, we find conditions such that the problem of finding an optimal device can be decoupled. Because the optimal device may still be difficult to compute, we also provide a simple algorithm that converges to a device that, although perhaps not optimal, still yields a ‘‘good’’ outcome for the designer. Our work in [19] demonstrates that these non-optimal protocols are very useful in a flow control management system.

APPENDIX A PROOF OF PROPOSITION 2

Proof: We prove \Rightarrow by contradiction.

Let $D = \langle (T_i), (A_i), \mu, X, \Phi \rangle$ be a maximum efficiency device (remember that we focus only on incentive compatible devices). Suppose **3** is not valid, i.e., there exists a type profile \tilde{t} such that the non-optimal action profile $z \neq g^M(\tilde{t})$ is suggested with positive probability $\mu_{\tilde{t}}(z) > 0$. Then

$$\begin{aligned} EU(f, g, D) &= \sum_{t \in T} \pi(t) \sum_{a \in A} \mu_t(a) \sum_{x \in X} \Phi_{t,a,a}(x) U(a, t, x) \\ &= V + W + \mu_{\tilde{t}}(z) \sum_{x \in X} \Phi_{\tilde{t},z,z}(x) U(z, \tilde{t}, x) \\ &< V + W + \mu_{\tilde{t}}(z) U(g^M(\tilde{t}), \tilde{t}, x^*) \leq EU^{ben} \end{aligned} \quad (4)$$

where

$$\begin{aligned} V &= \sum_{t \in T, t \neq \tilde{t}} \pi(t) \sum_{a \in A} \mu_t(a) \sum_{x \in X} \Phi_{t,a,a}(x) U(a, t, x) \\ W &= \pi(\tilde{t}) \sum_{a \in A, a \neq z} \mu_{\tilde{t}}(a) \sum_{x \in X} \Phi_{\tilde{t},a,a}(x) U(a, \tilde{t}, x) \end{aligned}$$

The first inequality of Eq. (4) is valid because of **A1** and **A2**, and the last inequality is valid because the benchmark optimum is the maximum achievable. Eq. (4) contradicts the fact that D is a maximum efficiency device.

Now suppose **4** is not valid, i.e., that $\Phi_a = \Phi_{t,g^M(t),a}$ does not sustain without intervention $g^M(t)$ in Γ_t . If Φ_a does not sustain $g^M(t)$ in Γ_t , then there exists a user i and an action

$a_i \neq g_i^M(t)$ such that user i prefers to adopt a_i when told to use $g_i^M(t)$, i.e., the strategy $g_i(t_i, g_i^M(t)) = a_i$ allows user i to obtain a higher utility with respect to the obedient strategy g_i^* ; this contradicts the fact that the device is incentive compatible. If Φ_a sustains $g^M(t)$ in Γ_t “with intervention”, then the users will adopt the action profile $g^*(t, g^M(t)) = g^M(t) \forall t$, but there exists \tilde{t} and $\tilde{x} \neq x^*$ such that $\Phi_{\tilde{t}, g^M(\tilde{t}), g^M(\tilde{t})}(\tilde{x}) > 0$. Then

$$\begin{aligned} EU(f, g^*, D) &= \sum_{t \in T} \pi(t) \sum_{x \in X} \Phi_{t, g^M(t), g^M(t)}(x) U(g^M(t), t, x) \\ &= V + W + \pi(\tilde{t}) \Phi_{\tilde{t}, g^M(\tilde{t}), g^M(\tilde{t})}(\tilde{x}) U(g^M(\tilde{t}), \tilde{t}, \tilde{x}) \\ &< V + W + \pi(\tilde{t}) \Phi_{\tilde{t}, g^M(\tilde{t}), g^M(\tilde{t})}(\tilde{x}) U(g^M(\tilde{t}), \tilde{t}, x^*) \leq EU^{ben} \end{aligned}$$

where

$$\begin{aligned} V &= \sum_{t \in T, t \neq \tilde{t}} \pi(t) \sum_{x \in X} \Phi_{t, g^M(t), g^M(t)}(x) U(g^M(t), t, x) \\ W &= \pi(\tilde{t}) \sum_{x \in X, x \neq \tilde{x}} \Phi_{\tilde{t}, g^M(\tilde{t}), g^M(\tilde{t})}(x) U(g^M(\tilde{t}), \tilde{t}, x) \end{aligned}$$

which contradicts the fact that D is a maximum efficiency device.

Finally, if **1** is not satisfied then **4** can not be satisfied either (because $g^M(t)$ is not sustainable without intervention), thus we obtain a contradiction. If **2** is not satisfied then either **3** is not satisfied or the device is not incentive compatible (because, given **3**, **2** is a particular case of the incentive-compatibility constraints), thus, in both cases, we obtain a contradiction.

⇐ Starting from the notion of equilibrium in Eq. (1), it is straightforward to verify that if **1** – **4** are satisfied the honest and obedient strategy profile (f^*, g^*) is sustained without intervention. Hence, the users always adopt the optimal action profile $g^M(t)$, the device uses the “non intervention action” x^* , and the utility of the designer is equal to the benchmark optimum defined in Eq. (2). ■

APPENDIX B PROOF OF LEMMA 2

Proof: Let $D = \langle (T_i), (A_i), \mu, X, \Phi \rangle$ be an optimal device (remember that we focus only on incentive compatible devices). The expected utility of user i having type t_i can be written as

$$\begin{aligned} EU_i(f, g, t_i, D) &= \sum_{t_{-i} \in T_{-i}} \pi(t | t_i) V_i(t) \\ V_i(t) &= \sum_{a \in A} \mu_t(a) \sum_{x \in X} \Phi_{t, a, a}(x) U_i(a, t, x) \end{aligned}$$

Denote by $\underline{a}_i(t)$ the minimum action recommended to user i with positive probability when the type profile is t , i.e., $\underline{a}_i(t) = \min \{a_i \in A_i : \mu_t(a_i, a_{-i}) > 0, a_{-i} \in A_{-i}\}$. We define the following intervals, $\forall i \in \{1, \dots, n\}$,

$$I_i(t) = \left[a_i^{min}, \min \left\{ \underline{a}_i(t), g_i^{NE^0}(t) \right\} \right]$$

and we use the notation $I(t)$ and $I_{-i}(t)$ in the usual way.

We define the function $\ell_i(a_{-i})$ in the domain $I_{-i}(t)$ as follows:

$$\ell_i(a_{-i}) = \{a_i \in I_i(t) \text{ such that } U_i(a, t, x^*) = V_i(t)\}$$

The function ℓ_i is a non-empty set-valued function from $I_{-i}(t)$ to the power set of $I_i(t)$. In fact, $\forall a'_{-i} \in I_{-i}(t)$,

$$\begin{aligned} U_i(a_i^{min}, a'_{-i}, t, x^*) &= 0 \leq V_i(t) \leq \sum_{a \in A} \mu_t(a) U_i(a, t, x^*) \\ &\leq \sum_{a \in A} \mu_t(a) U_i(a_i, a'_{-i}, t, x^*) \end{aligned} \quad (5)$$

The second inequality of Eq. (5) is valid because $U_i(a, t, x) \leq U_i(a, t, x^*)$, $\forall a, t, x$. The last inequality of Eq. (5) is valid because i 's utility is non increasing in the actions of the other users and, from the definition of the set $I_{-i}(t)$, $a'_{-i} \leq \underline{a}_{-i}(t) \leq a_{-i}$, $\forall a'_{-i} \in I_{-i}(t)$ and $\forall a_{-i}$ recommended with positive probability. Eq. (5) and the continuity of i 's utility imply that an action $\tilde{a}_i(t) \in I_i(t)$ satisfying $U_i(a, t, x^*) = V_i(t)$ exists, $\forall a_{-i} \in I_{-i}(t)$. Moreover, by definition $\ell_i(a_{-i})$ has a closed graph (i.e., the graph of $\ell_i(a_{-i})$ is a closed subset of $I(t)$) and, since i 's utility is non decreasing in $\left[a_i^{min}, g_i^{NE^0}(t) \right]$, $\ell_i(a_{-i})$ is convex, $\forall a_{-i} \in I_{-i}(t)$.

We define the function $\ell(a) = (\ell_1(a_{-1}), \dots, \ell_n(a_{-n}))$, $\forall a \in I(t)$. The function ℓ is defined from the non-empty, compact and convex set $I(t)$ to the power set of $I(t)$. Thanks to the properties of ℓ_i , ℓ has a closed graph and $\ell(a)$ is non-empty and convex. Therefore we can apply Kakutani fixed-point theorem [34] to affirm that a fixed point exists, i.e., there exists an action profile $\tilde{a}(t) \in I(t)$ such that $U_i(\tilde{a}, t, x^*) = V_i(t)$, $\forall i \in N$. Notice that $\tilde{a}(t) < g^{NE^0}(t)$, therefore $\tilde{a}(t)$ is sustainable without intervention in Γ_t , and we denote by Φ_a the intervention rule that sustains without intervention $\tilde{a}(t)$ in Γ_t .

Finally, the original optimal device $D = \langle (T_i), (A_i), \mu, X, \Phi \rangle$ can be substituted with the device $\tilde{D} = \langle (T_i), (A_i), \tilde{\mu}, X, \tilde{\Phi} \rangle$ in which, $\forall t$, $\tilde{\mu}(t) = \tilde{a}(t)$ and $\tilde{\Phi}_{t, \tilde{a}(t), a} = \Phi_a$. With the new device \tilde{D} the users are obedient (because the restriction of the intervention rule, Φ_a , sustains $\tilde{a}(t)$) and honest (because the utilities they obtain for each combination of reports are the same as in the initial device D that sustains the honest and obedient strategy profile). More specifically, \tilde{D} sustains without intervention the honest and obedient strategy profile. Moreover, in the equilibrium path the users' expected utilities using \tilde{D} coincide with the users' expected utilities using D ; thus, also the designer's utility (which is a function of users' utilities) remains the same, and this implies that \tilde{D} is optimal. ■

REFERENCES

- [1] G. Tan and J. Gutttag, “The 802.11 MAC protocol leads to inefficient equilibria,” in *Proc. 2005 IEEE INFOCOM*, vol. 1.
- [2] R. T. B. Ma, V. Misra, and D. Rubenstein, “Modeling and analysis of generalized slotted-Aloha MAC protocols in cooperative, competitive and adversarial environments,” in *Proc. 2006 IEEE ICDCS*.
- [3] J. Park and M. van der Schaar, “The theory of intervention games for resource sharing in wireless communications,” *IEEE J. Sel. Areas Commun.*, vol. 30, no. 1, pp. 165–175, Jan. 2012.
- [4] —, “Stackelberg contention games in multiuser networks,” *EURASIP J. Advances Signal Process.*, Jan. 2009.
- [5] —, “Designing incentive schemes based on intervention: the case of imperfect monitoring,” in *Proc. 2011 GameNets*.
- [6] Y. Xiao, J. Park, and M. van der Schaar, “Intervention in power control games with selfish users,” *IEEE J. Sel. Topics Signal Process.*, vol. 6, no. 2, pp. 165–179, Apr. 2012.
- [7] L. Hurwicz, “Game theory,” in *Decision and Organization: A Volume in Honor of Jacob Marshak*, 1972, pp. 297–336.

- [8] P. Dasgupta, P. Hammond, and E. Maskin, "The implementation of social choice rules: some results on incentive compatibility," *Review of Economic Studies*, vol. 46, no. 2, pp. 185–216, Apr. 1979.
- [9] B. Holmstrom, "Moral hazard and observability," *Bell J. Economics*, vol. 10, pp. 74–91, Aug. 1979.
- [10] R. B. Myerson, "Incentive-compatibility and the bargaining problem," *Econometrica*, vol. 47, no. 1, pp. 61–73, Jan. 1979.
- [11] —, "Optimal auction design," *Mathematics of Operations Research*, vol. 6, pp. 58–73, Dec. 1981.
- [12] —, "Optimal coordination mechanism in generalized principal-agent problems," *J. Mathematical Economics*, vol. 10, no. 1, pp. 67–81, Jun. 1982.
- [13] J. Huang, R. A. Berry, and M. L. Honig, "Distributed interference compensation for wireless networks," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 5, pp. 1074–1084, May 2006.
- [14] Y. Su and M. van der Schaar, "Linearly coupled communication games," *IEEE Trans. Commun.*, vol. 59, no. 9, pp. 2543–2553, Sep. 2011.
- [15] Y. Gai, H. Liu, and B. Krishnamachari, "A packet dropping-based incentive mechanism for M/M/1 queues with selfish users," in *Proc. 2011 IEEE INFOCOM*, pp. 2687–2695.
- [16] K. Bharath-Kumar and J. M. Jaffe, "A new approach to performance-oriented flow control," *IEEE Trans. Commun.*, vol. 29, no. 4, pp. 427–435, Apr. 1981.
- [17] C. Douligeris and R. Mazumdar, "A game theoretic perspective to flow control in telecommunication networks," *J. Franklin Institute*, vol. 329, no. 2, pp. 383–402, Mar. 1992.
- [18] Z. Zhang and C. Douligeris, "Convergence of synchronous and asynchronous greedy algorithm in a multiclass telecommunications environment," *IEEE Trans. Commun.*, vol. 40, no. 8, pp. 1277–1281, Aug. 1992.
- [19] L. Canzian, Y. Xiao, W. Zame, M. Zorzi, and M. van der Schaar, "Intervention with complete and incomplete information: application to flow control," to appear in *IEEE Trans. Commun.*, 2013.
- [20] L. Canzian, L. Badia, and M. Zorzi, "Promoting cooperation in wireless relay networks through Stackelberg dynamic scheduling," *IEEE Trans. Commun.*, vol. 61, no. 2, 2013, pp. 700–711.
- [21] D. Wang, C. Comaniciu, and U. Tureli, "Cooperation and fairness for slotted Aloha," *Wireless Personal Commun.*, vol. 43, no. 1, pp. 13–27, Oct. 2007.
- [22] L. Yang, H. Kim, J. Zhang, M. Chiang, and C. W. Tan, "Pricing-based spectrum access control in cognitive radio networks with random access," in *Proc. 2011 IEEE INFOCOM*, pp. 2228–2236.
- [23] T. Basar and R. Srikant, "Revenue-maximizing pricing and capacity expansion in a many-users regime," in *Proc. 2002 IEEE INFOCOM*, pp. 23–27.
- [24] H. Shen and T. Basar, "Optimal nonlinear pricing for a monopolistic network service provider with complete and incomplete information," *IEEE J. Sel. Areas Commun.*, vol. 25, no. 6, pp. 1216–1223, 2007.
- [25] D. A. Schmidt, C. Shi, R. A. Berry, M. L. Honig, and W. Utschick, "Distributed resource allocation schemes," *IEEE Signal Process. Mag.*, vol. 26, no. 5, pp. 53–63, Sep. 2009.
- [26] T. Alpcan, T. Basar, R. Srikant, and E. Altman, "CDMA uplink power control as a noncooperative game," *Wireless Netw.*, vol. 8, no. 6, pp. 659–670, Nov. 2002.
- [27] J. Huang, R. A. Berry, and M. L. Honig, "Auction-based spectrum sharing," *Mobile Networks and Applications*, vol. 11, no. 3, pp. 405–418, Jun. 2006.
- [28] T. Alpcan, H. Boche, and S. Naik, "A unified mechanism design framework for networked systems," *CoRR*, vol. abs/1009.0377, 2010.
- [29] P. Nuggehalli, J. Price, and T. Javidi, "Pricing and QoS in wireless random access networks," in *Proc. 2008 IEEE GLOBECOM*.
- [30] W. Vickrey, "Counterspeculation, auctions, and competitive sealed tenders," *J. Finance*, vol. 16, no. 1, pp. 8–37, 1961.
- [31] J. Park and M. van der Schaar, "A note on the intervention framework," UCLA University, Electrical Engineering Department, Tech. Rep., 2011. Available: http://medianetlab.ee.ucla.edu/papers/note_intervention.pdf
- [32] G. Owen, *Game Theory*, 3rd ed. Academic, 2001.
- [33] R. J. Aumann, "Subjectivity and correlation in randomized strategies," *J. Mathematical Economics*, vol. 1, no. 1, pp. 67–96, Mar. 1974.
- [34] S. Kakutani, "A generalization of Brouwer's fixed point theorem," *Duke Mathematical J.*, vol. 8, no. 3, pp. 457–459, 1941.

Luca Canzian [M'13] received the B.Sc., M.Sc., and Ph.D. degrees in Electrical Engineering from the University of Padova, Italy, in 2005, 2007, and 2013, respectively. From 2007 to 2009 he worked in Venice, Italy, as an R&D Engineer at Tecnomare, a company providing design and engineering services for the oil industry. From September 2011 to March 2012 he was on leave at the University of California, Los Angeles (UCLA). Since January 2013, he has been a PostDoc at the Electrical Engineering Department at UCLA. His research interests include resource allocation, game theory, online learning and real-time stream mining.

Yuanzhang Xiao received the B.E. and M.E. degree in Electrical Engineering from Tsinghua University, Beijing, China, in 2006 and 2009, respectively. He is currently pursuing the Ph.D. degree in the Electrical Engineering Department at the University of California, Los Angeles. His research interests include game theory, optimization, communication networks, and network economics.

William Zame (Ph.D., Mathematics, Tulane University 1970) is Distinguished Professor of Economics and Mathematics at UCLA. He previously held appointments at Rice University, SUNY/Buffalo, Tulane University and Johns Hopkins University, and visiting appointments at the Institute for Advanced Study, the University of Washington, the Institute for Mathematics and its Applications, the Mathematical Sciences Research Institute, the Institut Mittag-Leffler, the University of Copenhagen, VPI, UC Berkeley and Einaudi Institute for Economics and Finance. He is a Fellow of the Econometric Society (elected 1994) and a former Guggenheim Fellow (2004-2005).

Michele Zorzi [F'07] received the Laurea and the PhD degrees in Electrical Engineering from the University of Padova, Italy, in 1990 and 1994, respectively. During the Academic Year 1992/93, he was on leave at the University of California, San Diego (UCSD) as a visiting PhD student, working on multiple access in mobile radio networks. In 1993, he joined the faculty of the Dipartimento di Elettronica e Informazione, Politecnico di Milano, Italy. After spending three years with the Center for Wireless Communications at UCSD, in 1998 he joined the School of Engineering of the University of Ferrara, Italy, where he became a Professor in 2000. Since November 2003, he has been on the faculty at the Information Engineering Department of the University of Padova. His present research interests include performance evaluation in mobile communications systems, random access in mobile radio networks, ad hoc and sensor networks, energy constrained communications protocols, broadband wireless access and underwater acoustic communications and networking.

Dr. Zorzi was the Editor-In-Chief of the *IEEE Wireless Communications Magazine* from 2003 to 2005 and the Editor-In-Chief of the *IEEE TRANSACTIONS ON COMMUNICATIONS* from 2008 to 2011, and currently serves on the Editorial Board of the *Wiley Journal of Wireless Communications and Mobile Computing*. He was also guest editor for special issues in the *IEEE Personal Communications Magazine* (Energy Management in Personal Communications Systems), *IEEE Wireless Communications Magazine* (Cognitive Wireless Networks), and the *IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS* (Multi-media Network Radios, and Underwater Wireless Communications Networks). He served as a Member-at-large of the Board of Governors of the IEEE Communications Society from 2009 to 2011.

Mihaela van der Schaar [F'10] is Chancellor's Professor of Electrical Engineering at University of California, Los Angeles. Her research interests include network economics and game theory, online learning, multimedia networking, communication, processing, and systems, real-time stream mining, dynamic multi-user networks and system designs. She is an IEEE Fellow, a Distinguished Lecturer of the Communications Society for 2011-2012, the Editor in Chief of *IEEE TRANSACTIONS ON MULTIMEDIA* and a member of the Editorial Board of the *IEEE JOURNAL ON SELECTED TOPICS IN SIGNAL PROCESSING*. She received an NSF CAREER Award (2004), the Best Paper Award from *IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY* (2005), the Okawa Foundation Award (2006), the IBM Faculty Award (2005, 2007, 2008), the Most Cited Paper Award from *EURASIP: Image Communications Journal* (2006), the Gamenets Conference Best Paper Award (2011) and the 2011 IEEE Circuits and Systems Society Darlington Award Best Paper Award. She received three ISO awards for her contributions to the MPEG video compression and streaming international standardization activities, and holds 33 granted US patents. For more information about her research visit: <http://medianetlab.ee.ucla.edu/>