Efficient Allocations in Economies with Asymmetric Information when the Realized Frequency of Types is Common Knowledge

Aristotelis Boukouras, University of Leicester

Kostas Koufopoulos, University of Pireaus

Efficient Allocations in Economies with Asymmetric Information when the Realized Frequency of Types is Common Knowledge

Aristotelis Boukouras[*] Kostas Koufopoulos[†]

March 31, 2015

Abstract

We consider a general economy, where agents have private information about their types. Types can be multi-dimensional and potentially interdependent. We show that, if the realized frequency of types (the exact number of agents for each type) is common knowledge, then a mechanism exists, which is consistent with truthful revelation of private information and which implements first-best allocations of resources as the unique equilibrium. The result requires the single crossing property on utility functions and the anonymity of the Pareto correspondence.

Keywords: adverse selection, first-best, full implementation, mechanism design, single-crossing property

JEL Classification: D71, D82, D86

[*]University of Leicester, University Road 1, LE1 7RH, Leicester, UK, email: Aris.Boukouras@le.ac.uk

[†]University of Pireaus, Department of Banking and Financial Management, Pireaus, Greece, koufop@unipi.gr

1

1 Introduction

As first shown by Akerlof (1970), Spence (1973) and Rothschild and Stiglitz (1976), hidden-types (adverse selection) problems can have significant consequences in terms of efficiency on economic outcomes. More specifically, incentive compatibility constraints limit the set of feasible allocations that can be attained. How are these restrictions relaxed as more information becomes common knowledge? And what is the minimum additional information required for achieving first-best efficiency? These are some of the questions that have emerged in the attempt to better understand the effects of information aggregation on efficiency. Indeed, some early papers by McAfee (1992), Armstrong (1999) and Casella (2002) already point toward this direction.

In this paper, we claim that if the number of agents with the same type is known for all types in a population (in other words, the realized frequency of types is known), then it is possible, under general conditions, to implement first-best allocations as a unique equilibrium. More precisely, we consider an economy with asymmetric information, where each agent has private information about his type. We also assume that: (i) the realized frequency of types is common knowledge, (ii) preferences satisfy the single crossing property, and (iii) the social choice rule satisfies anonymity. Given these conditions, we show that it is possible to construct a mechanism which has a unique equilibrium, where all agents reveal their type truthfully and they receive a first-best allocation.

The result is interesting because we examine an asymmetric information problem which is situated in-between the problem of Maskin (1999) (in which all agents know the state of world but the mechanism designer does not know it) and the classic adverse selection (in which each agent knows only his own type and the mechanism designer knows the ex-ante distribution of types). The intuition behind the result is that, if the realized frequency of types is known, then one can aggregate the messages that all agents are sending out and uncover any misreport(s), even if the identity of the liar is not known. That is, appropriately designed punishments for lying can induce agents to reveal their information truthfully. We talk about appropriately designed punishments, because one of the features of our mechanism is that punishments must not be too harsh. If the punishment when a lie is detected is too severe, then some agents may deliberately lie about their type in order to force other agents to also do so. The lies cancel out at the aggregate level and the former agents "steal" the allocations of the latter, who are forced to lie under the fear of the extreme punishments. This can

lead to coordination failures and multiplicity of equilibria. Therefore, uniqueness of the equilibrium requires a careful construction of the off-the-equilibrium path allocations when lies are detected. We show that such punishments exist when the indifference curves satisfy the single crossing property.

This result is also interesting for two more reasons. First, one may consider economic applications with a finite number of agents, where, in addition to the private information that each individual has, there is knowledge about how many agents have each type. This additional information could come from a positive or negative informational shock. For example, a retail store has received pre-paid orders from its customers, has already the goods in stock and is ready to make the deliveries. However, the records on the orders get destroyed due to an accident and the store's manager does not know who made each order. What is he to do? Can he induce the customers to truthfully reveal the orders they have made without them making unreasonable claims or receiving orders that were meant for other customers? We claim that this is possible, as long as the manager posts a list with all the orders made and gives to each customer a basket of goods, which depends on how many other agents have claimed to have ordered it.

Second, there are some well-known models of adverse selection (for example Akerloff (1970) and Spence (1973)) which assume that the proportion of each type in the population is common knowledge. For these models, the mechanism presented in this paper can be used in order to provide first-best allocations. To the best of our knowledge, this efficiency result has not been provided in the literature so far.

The most closely related paper to ours is Jackson and Sonnenschein (2007), who consider an economy where agents play multiple copies of the same game at the same time and their types are independently distributed across games. They allow for mechanisms, which "budget" the number of times that an agent claims to be of a certain type. If the number of parallel games becomes very large, then all the Bayes-Nash equilibria of these mechanisms converge to first-best allocations. Our model is different from theirs, because we do not require multiple games to be played at the same time but we impose a stronger assumption on what is common knowledge. Moreover, we allow for interdependent values, while they consider an independent values setting, and in our model asymmetric information may include other individual characteristics apart from preferences (productivity parameters, proneness to accidents, etc.).

McLean and Postlewaite (2002, 2004) consider efficient mechanisms in economies with interdependent values. The state of the world is unknown to all agents, but each individual receives a noisy private signal about the state. They show that when

signals are sufficiently correlated with the state of the world and each agent has small informational size (in the sense that his signal does not contain additional information about the state of the world when the signals of all the other agents are taken into account), then their mechanism implements allocations arbitrarily close to first-best allocations. However, in the model of McLean and Postlewaite, when private signals are perfectly correlated, all agents learn not only their own type but also the type of all other agents. That is, in the limit, the framework of McLean and Postlewaite is one of complete information. In contrast, in our setting agents know, at most, the realized frequency of types. [1]

VCG-mechanisms (Vikrey, 1961, Clarke, 1971, Groves, 1973) are often reference points in terms of results on efficiency. With respect to these mechanisms, our paper is more general as they assume quasi-linear preferences while we allow for guasi-concave utility functions. Moreover, these papers show that the respective mechanisms that they examine produce truth-telling equilibria, but they do not examine whether other equilibria, non-truth-telling, exist. In contrast, we consider this possibility and show that the truth-telling equilibrium of our mechanism is unique.

Our paper is also related to the auctions literature with interdependent types. In this context, Crémer and McLean (1985) and Perry and Reny (2002, 2005), show the existence of efficient auctions when types are interdependent. Crémer and McLean, however, require quasi-linear preferences while we do not. Perry and Reny are closer to our result since they also assume that the single crossing property holds. Nonetheless, our main focus is the uniqueness of the equilibrium, an issue which, as with the VCG literature, is not studied in these papers.

Rustichini, Satterthwaite and Williams (1994) show that the inefficiency of trade between buyers and sellers of a good, who are privately informed about their preferences, rapidly decreases with the number of agents involved in the two sides of the market and in the limit it reaches zero. Effectively, the paper examines the issue of convergence to the competitive equilibrium as the number of agents increases. However, their model is limited to private values problems and hence it can be seen as a special case of our formulation.

More recently, the papers by Mezzetti (2004) and Ausubel (2004),(2006) examine the issues of efficient implementation under interdependent valuations and independently distributed types. However, they also assume that agents' preferences are quasi-linear

---

[1]In a sense, in our model agents receive private signals as well, but one can think of them as perfect signals about the frequency of types.

4

with respect to the transfers they receive, whereas in our model utility may not be transferable. Moreover, the mechanisms proposed in these papers may generate multiple equilibria (in most of which truth-telling is violated), while we are interested in a mechanism which has a unique truth-telling equilibrium.

Finally, several recent papers examine efficient mechanism design in dynamic settings. The most notable papers in this category are the papers by Battaglini (2005), Athey and Segal (2007), Gershkov and Moldovanu (2009), Bergemann and Välimäki (2010), Pavan et al (2014), Athey and Segal (2013) and Escobar and Toikka (2013). Our paper differs significantly from these papers. We assume that types are drawn only once and the realized distribution of types becomes common knowledge subsequently, while they assume that agents' type evolves over time according to a stochastic process which is common knowledge. On the other hand, we also use a multi-stage mechanism in order to induce truthful reporting. As a result, even though agents' private information does not change in the various stages, their incentive to report truthfully changes according to the information they learn from the previous stages, similarly to the dynamic mechanism design literature.

2 An Example: Spence (1973)

First we demonstrate how the knowledge of the realized frequency of types can be used to implement first best allocations as a unique equilibrium by applying the main idea to the classic paper by Spence (1973). The economy consists of two types of workers. Type 1 has low productivity $\underline{a}$ and its proportion of the population is q_1. Type 2 has high productivity $\overline{a}$, $(\overline{a} > \underline{a})$ and its proportion of the population is $1 - q_1$.[2] Acquiring y units of education costs $y/\underline{a}$ for type 1 and $y/\overline{a}$ for type 2. Productivity parameters are private information and firms hire workers according to a wage schedule, based on verifiable educational attainment. The payoff for an individual is the value of his wage minus the educational cost and for a firm the productivity parameter minus the wage.

Spence argues that agents will acquire education (which does not increase productivity in his model) in order to signal their productivity to firms. In equilibrium, the wage schedules are such that high productivity workers acquire some education and

---

[2]Note that in the original paper, Spence made the assumption that a known proportion of the population belongs to one type and the remainder proportion belongs to the other type. Hence, he implicitly made the assumption that the realized frequency of types is common knowledge and, hence, we can apply our mechanism directly into his economy.

credibly signal their type, while low productivity workers acquire no education, and firms correctly infer that they are of low productivity. The education acquired by type 2 is a deadweight loss, but necessary for credible signaling.
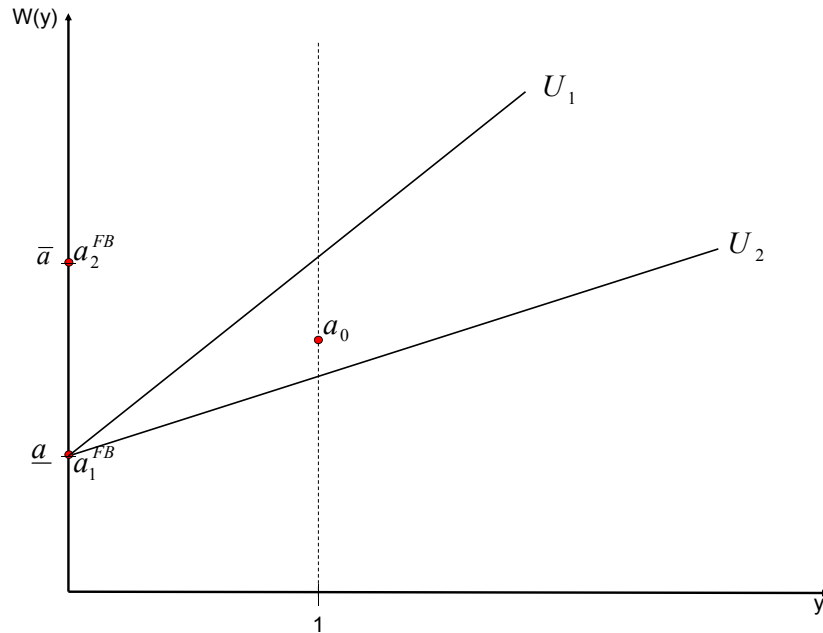

Figure 1: Spence, 1973

Assume that the total population is N. Then Nq_1 is the total number of agents of type 1 and $N(1 - q_1)$ is the total number of agents of type 2. Given this, the following mechanism can separate types without any agent incurring educational costs in equilibrium. Let all workers report their type. If the number of agents who report type 1 and 2 is Nq_1 and $N(1 - q_1)$ respectively, then agents who report type 1 receive wage $w_1 = \underline{a}$ and zero education (contract α_1^{FB} in figure 1) and those who report type 2 receive wage $w_2 = \overline{a}$ and zero education (contract α_2^{FB} in figure 1). In any other case, where the reported number of types do not match their population size, those who report type 1 receive $w_1 = \underline{a}$ and those who report type 2, are asked to undertake one unit of education and receive $w_2 = \underline{a} + \epsilon$, with $\frac{1}{\overline{a}} < \epsilon < \frac{1}{\underline{a}}$ (recall that a unit of education costs $\frac{1}{\overline{a}}$ for high productivity workers and $\frac{1}{\underline{a}}$ for low productivity workers).

The above mechanism fully implements the first-best allocations in this economy. First, consider the strategies of type 2. It is clear that, irrespectively of the reports of the other agents, it is a strictly dominant strategy for him to report his type truthfully. This

6

is because, when everybody else reports truthfully, type 2 prefers to report truthfully as well (then his payoff is $\bar{a}$) than to misreport his type (then his payoff is $\underline{a}$), given that $\bar{a} > \underline{a}$. Similarly, if someone else lies, type 2 prefers to report truthfully and receive a payoff of $\underline{a} + \epsilon - \frac{1}{\bar{a}}$ than to cover the lie by misreporting and receive $\underline{a}$, given that $\underline{a} + \epsilon - \frac{1}{\bar{a}} > \underline{a}$. Given the dominant strategy of type 2 and $\underline{a} > \underline{a} + \epsilon - \frac{1}{\underline{a}}$, it is a best-response for type 1 to report truthfully as well. Hence, all agents report truthfully in equilibrium and acquire zero education. In Figure 1, contract a_0 denotes the offer to the workers, who report high productivity, when lies are detected.

3   The Economy

The previous example was used in order to show that it is possible to eliminate asymmetric information problems if the realized frequency is common knowledge. We now proceed to show that this result is general and does not depend on the specifics of the example. First, we introduce the economy and the notation.

The economy consists of a finite set I of agents, with I standing for the aggregate number of agents as well. Θ is the finite set of potential types with elements ϑ. In order to make our problem non-trivial we assume that $I \geq 2$ and $\Theta \geq 2$. Each agent has private information about his own type, but does not know the types of the other agents. β is the vector of realized frequencies of types in the population. That is β denotes the *ex post* distribution of types in the population, i.e. the relative frequency of each type, which materializes after types are drawn. Therefore, $\beta(\vartheta)$ is the proportion of agents who have type ϑ in the population and $N(\vartheta)$ is the total number of agents of type ϑ: $N(\vartheta) = \beta(\vartheta)I$.

Let A be the set of all feasible allocations, with elements $a \in A \subseteq \boldsymbol{R}_+^{I \times L}$, with $L \geq 2$. L can be interpreted as the number of commodities in the economy. Also, for any subset J of the set I, let A^J be the set of feasible allocations for the agents in J ($A^J \subseteq \boldsymbol{R}_+^{J \times L}$). For the analysis that follows it is also useful to define allocations on an individual basis. That is, given an allocation $a \in A$, the **individual allocation** $a_i \in \boldsymbol{R}_+^L$ is the bundle that agent i consumes. Moreover, since later on we will require that agents of the same type consume the same bundle, it is useful to denote individual allocations with respect to types. That is, a_ϑ denotes the individual allocation that an agent of type ϑ consumes within allocation a.

$u : \boldsymbol{R}^L \times \Theta \times \Theta^{I \setminus i} \to \boldsymbol{R}$ is the *Bernoulli* utility function for agent i, which we

assume to be strictly quasi-concave. The following definitions are also useful. $L_\vartheta(a_\vartheta)$ is the **lower-contour set** of an agent with type ϑ associated with individual allocation a_ϑ: $L_\vartheta(a_\vartheta) = \{c \in \mathbf{R}_+^L : u_\vartheta(c) < u_\vartheta(a_\vartheta)\}$. $V_\vartheta(a_\vartheta)$ is the **upper-contour set** of type ϑ associated with a_ϑ: $V_\vartheta(a_\vartheta) = \{c \in \mathbf{R}_+^L : u_\vartheta(c) > u_\vartheta(a_\vartheta)\}$.

Overall, the economy is described by the following primitives: $E = \{I, A, u, \Theta, \beta\}$. This formulation of the economy allows for modeling a wide variety of economic situations. Since we impose no restrictions on β or the type-generating process that produces β, types may or may not be independently distributed. Moreover, the utility function of agents may or may not depend on the types of other agents, and so both adverse-selection problems with independent or inter-dependent valuations can be seen as special cases of our formulation. The model also allows for public goods problems, since some elements of the individual allocations can be common.

Economies with uncertainty can be easily accommodated by our model as well. For example, let $\phi : \beta \to \Delta^S$ be the probability distribution function over states, where S the finite set of states and Δ^S is the unit simplex $\{\phi \in \mathbf{R}_+^S | \sum_{s \in S} \phi_s = 1\}$. In this case, $L = S \times T$, where T is the finite set of final commodities, and the agents' expected utility function is $u_i(a_i, \beta) = \sum_{s \in S} v_i(a_i, s)\phi_s(\beta)$, where $v_i(a_i, s)$ is the decision-outcome payoff in state s.

4 Implementation of First Best Allocations

In this section we provide the main result of the paper. It is shown that if preferences satisfy the single-crossing condition and the first-best allocation is *anonymous* (i.e. individual allocations depend on agents' type and not on their identity), then there exists a mechanism which implements the first-best allocation as a unique equilibrium.

In order to prove this result we proceed as follows. First, we define *anonymity* in our setting (Definition 1 and Assumption 1) and we assume that types' preferences satisfy the single-crossing condition (Definition 2 and Assumption 2). We then show that Pareto efficiency implies a ranking of types according to envy (Lemma 1) which is exploited in the design of the out-of-equilibrium path allocations. The combination of the single-crossing condition with the result of Lemma 1 allows one to construct *incentive compatible* individual allocations, according to the notion of incentive compatibility provided by Definition 3, for any first-best individual allocation (Lemma 2). Finally, Lemma 1 and 2 are combined for the main result (Theorem 1), in which it is shown

that there exists a mechanism which fully implements any first-best allocation.

Formally, let $\boldsymbol{a}^*$ be a Pareto efficient (first-best) allocation and let a^*_ϑ be the individual allocation, which an agent with type ϑ receives in $\boldsymbol{a}^*$. In other words, a^*_ϑ is the individual allocation which a mechanism designer would like to offer to an agent with type ϑ, if $\boldsymbol{a}^*$ were to be implemented. Then, we have the following definition:

Definition 1: A Social Choice Rule satisfies Anonymity if, for any two agents i and j, $a^*_i = a^*_j = a^*_\vartheta$ whenever $\vartheta_i = \vartheta_j = \vartheta$.

Assumption 1: The Social Choice Rule satisfies Anonymity.

Under Anonymity, agents who have identical types receive identical individual allocations. Therefore, an agent's identity per-se has no impact on the agent's final individual allocation. Anonymity is a desirable property for a social choice rule. In most cases of interest, economists are concerned with the economic characteristics of agents and not with their identity. Therefore, it is reasonable to assume that, if the distribution of these characteristics remains unchanged, so does the distribution of the economically desirable outcomes. It is also a property satisfied by many commonly used social choice rules, like the Walrasian correspondence and the utilitarian social welfare function.

Definition 2 (Single-Crossing Condition) For any two types ϑ and η, $\{\vartheta, \eta\} \in \Theta$, there exists at least one pair of commodities $k, l \in L$ such that $-\frac{\partial u_\vartheta / \partial l}{\partial u_\vartheta / \partial k} < -\frac{\partial u_\eta / \partial l}{\partial u_\eta / \partial k}$.

Assumption 2: Preferences for all types satisfy the single-crossing condition.

We now show that any Pareto efficient allocation $\boldsymbol{a}^*$ implies a ranking of types according to envy.

Lemma 1: If $\boldsymbol{a}^*$ is a Pareto efficient allocation which satisfies Anonymity, then there exists at least one type ϑ, who does not envy the individual allocation of any other agent: $U_\vartheta(a^*_\vartheta) \geq U_\vartheta(a^*_\eta), \forall \eta \in \Theta$.

Proof: Consider a Pareto efficient allocation $\boldsymbol{a}^*$, which satisfies Anonymity with type-dependent individual allocations a^*_ϑ and suppose that Lemma 1 does not hold. Then, all types envy at least one other type: $\forall a^*_\vartheta, \ \exists \eta \in \Theta, \eta \neq \vartheta : U_\vartheta(a^*_\eta) > U_\vartheta(a^*_\vartheta)$. But,

since this holds for all types, then there exists at least one reassignment of individual allocations among the I individuals such that some of them are made strictly better-off and the rest remain as well-off as under $\boldsymbol{a}^*$.

In order to find one such reassignment, use the following algorithm. Pick an arbitrary $\vartheta \in \Theta$ and let $\overline{\vartheta} = \{\eta \in \Theta : U_\vartheta(a_\eta^*) > U_\vartheta(a_\vartheta^*)\}$, be the set of types whom ϑ-types envy. Reassign one individual allocation a_η^*, for some $\eta \in \overline{\vartheta}$, to one agent of type ϑ. If $\vartheta \in \overline{\eta}$, then reassign a_ϑ^* (from the ϑ-type individual who received a_η^*) to η (to the specific agent whose a_η^* individual allocation was reassigned) and stop the reassignment.

If $\vartheta \notin \overline{\eta}$, then reassign some individual allocation a_ζ^*, $\zeta \in \overline{\eta}$ to η and then proceed to the individual whose individual allocation a_ζ^* was reassigned. Iterate the procedure until you reach some agent of type λ, such that there exists some type $\kappa \in \overline{\lambda}$, whose individual allocation a_κ^* has already being reassigned. In this case, ignore all reassignments preceding the individual of type κ (these agents retain their original individual allocations), reassign to λ the individual allocation a_κ^* and stop the reassignments (all reassignments between κ and λ are not modified).

Since the set of agents is finite and all types envy at least one individual allocation, after at most I reassignments, the algorithm above will end up in some agent, whose individual allocation has already been reassigned. In this case, a reassignment of individual allocations has been found, which makes some agents in I better-off (from agent of type κ until agent λ) while the rest remain equally well-off. This constitutes a Pareto improvement and violates the initial assumption that $\boldsymbol{a}^*$ is Pareto efficient. ∎

Corollary 1: If $\boldsymbol{a}^*$ is a Pareto efficient allocation which satisfies Anonymity, then Lemma 1 holds for any subset of Θ. Namely, let $\check{\Theta} \subseteq \Theta$ and let $\check{A} = \{a_\vartheta^* : \vartheta \in \check{\Theta}\}$. Then, Lemma 1 holds for $\check{\Theta}$ with regard to $\check{A}$ as well.

Proof: Take any subset of agents $\check{\Theta}$ of the set Θ. Suppose that Lemma 1 does not hold over the set $\check{A}$, which is the set of individual allocations of the agents with types in $\check{\Theta}$. But if Lemma 1 does not hold, then it is possible to find a reassignment of allocations between the agents in $\check{\Theta}$, such that some of them will be made better-off while the rest remain as well-off. This is a Pareto-improvement for some agents in I, which contradicts the assumption that $\boldsymbol{a}^*$ is Pareto efficient. ∎

Lemma 1 and Corollary 1 allows us to construct a complete ranking of types according

to envy. To see this, let $K = \{\vartheta \in \Theta : U_\vartheta(a_\vartheta^*) \geq U_\eta(a_\eta^*), \forall \eta \in \Theta\}$ be the set of types who do not envy the individual allocation of any other type. By Lemma 1, we know that this set is non-empty. Then, by removing this set of types from the set Θ and applying Corollary 1, we can define $K - 1 = \{\vartheta \in \Theta : U_\vartheta(a_\vartheta^*) \geq U_\eta(a_\eta^*), \forall \eta \in \Theta - K\}$. By iteration, we can define K envy groups, $1 \leq K \leq \Theta$, such that the types in each one of them (say envy group k) do not envy any of the types in their own group or lower groups (any $l < k$), but they envy some type(s) in higher groups (some type in an envy group $h > k$).[3]

The K envy groups defined above could in principle be used to construct a mechanism which induces all types to reveal their type truthfully, however, such a mechanism would involve tedious case distinctions across possible groups. In order to make the workings of the mechanism more transparent we are going to rank types within each envy group that contains multiple types so that the mechanism is implemented over a full ranking of types. Since the way these envy groups are constructed ensures that there is no envy between types which belong to the same group, one can rank types within envy groups to get a complete ranking of types. In what follows we assume that the mechanism designer ranks types according to the following simple rules:

1. Types who belong to a higher envy group are ranked above types who belong to a lower envy group.

2. If two types, ϑ and η, belong to the same envy group and $u_\vartheta(a_\vartheta^*) > u_\vartheta(a_\eta^*)$, $u_\eta(a_\eta^*) = u_\eta(a_\vartheta^*)$, then type ϑ receives higher ranking than type η.

3. If two types, ϑ and η, belong to the same envy group and $u_\vartheta(a_\vartheta^*) > u_\vartheta(a_\eta^*)$, $u_\eta(a_\eta^*) > u_\eta(a_\vartheta^*)$, then the ranking between the two types is arbitrarily determined as long as it is compatible with rules 1. and 2. above whenever comparing the rank of types ϑ and η with the rest of the types.

Note that, in principle the case where ϑ and η belong to the same envy group and $u_\vartheta(a_\vartheta^*) = u_\vartheta(a_\eta^*)$, $u_\eta(a_\eta^*) = u_\eta(a_\vartheta^*)$ should be examined. However, this case is incompatible with Assumption 2 on single crossing and so it is omitted from the above list. Overall, by following the above rules, the mechanism designer ranks all types

---

[3]One extreme case is when an allocation exhibits no-envy, in which case K contains the whole set of types (egalitarian allocations: $K = 1$) The other extreme case is when each envy group contains a single type, in which case the types form a complete hierarchy, from the one who is envied by all the other types to the one who is not envied by anyone ($K = \Theta$).

according to envy from the lowest one, type 1, to the highest, type Θ. From this point forward we use ϑ to denote the rank of a type, so that $1 \leq \vartheta \leq \Theta$. By construction of this ranking, type Θ does not envy the first-best individual allocation of any other type. A generic type ϑ may envy the first-best individual allocation of a higher type ($\kappa > \vartheta$), but does not envy the first-best individual allocation of any type with lower rank ($\eta < \vartheta$). For the rest of the paper, we refer to type Θ as the type with the **highest rank** and type 1 as the type with the **lowest rank**.[4]

In Lemma 2 below we exploit the ranking of types and the single crossing property to show that it is possible to contract incentive compatible individual allocations for a subset of types given an arbitrary first-best individual allocation as reference point. Before providing this result, however, we need to clarify the notion of incentive compatibility used in this paper, which is done by Definition 3.

Definition 3: Given a type η and his respective first-best individual allocation a_η^* as a reference point, the set $\hat{A}(a_\eta^*)$ of individual allocations is *incentive compatible* if $a_\eta^* \in \hat{A}(a_\eta^*)$ and for any type ϑ, $\eta < \vartheta \leq \Theta$, there exists an individual allocation $\hat{a}_\vartheta(a_\eta^*) \in \hat{A}(a_\eta^*)$ such that $u_\vartheta(\hat{a}_\vartheta(a_\eta^*)) > u_\vartheta(\hat{a}_\kappa(a_\eta^*))$, $\forall$ $\hat{a}_\kappa(a_\eta^*) \in \hat{A}(a_\eta^*)$, $\kappa \neq \vartheta$.

In words, given a subset of types in which η is the lowest rank, definition 3 requires that a set of individual allocations is incentive compatible if for each type with rank higher than η there is an individual allocation which he prefers over all other individual allocations in the set. Lemma 2 below shows that it is always possible to construct a set $\hat{A}(a_\eta^*)$ which satisfies the definition above.

Lemma 2: If Assumptions 1 and 2 hold then for any type $\eta \in \Theta$, with $\eta < \Theta$, there exists a set $\hat{A}(a_\eta^*)$, which is incentive compatible.

Proof: Take $\boldsymbol{a}^*$ as given and consider an arbitrary type η ($\eta < \Theta$) with corresponding first-best individual allocation a_η^*. Also recall that, by Lemma 1 and Corollary 1, there is a complete ranking of types from 1 to Θ. We use this ranking and definition 2 to construct a sequence of individual allocations $\hat{a}_\eta(a_\eta^*)$ to $\hat{a}_\Theta(a_\eta^*)$ which satisfy the definition of incentive compatibility in definition 3.

---

[4]Note that the ranking of types depends implicitly on the implementable allocation $\boldsymbol{a}^*$. Different allocations will lead to different rankings. But since $\boldsymbol{a}^*$ is common knowledge, the ranking of types is common knowledge as well.

To prove this result, we consider the following algorithm for picking a sequence of such individual allocations. Take a_η^* as a starting point $(\hat{a}_\eta(a_\eta^*) \equiv a_\eta^*)$ and pass the indifference curves of types η and $\eta + 1$ from this point. By Definition 2, the set $V_{\eta+1}(a_\eta^*) \cap L_\eta(a_\eta^*)$ is non-empty. Pick an individual allocation in the set $V_{\eta+1}(a_\eta^*) \cap L_\eta(a_\eta^*)$ as $\hat{a}_{\eta+1}(a_\eta^*)$ (possibly arbitrarily close to a_η^*). Note that, by construction, η strictly prefers $\hat{a}_\eta(a_\eta^*)$ to $\hat{a}_{\eta+1}(a_\eta^*)$, while $\eta + 1$ strictly prefers $\hat{a}_{\eta+1}(a_\eta^*)$ to $\hat{a}_\eta(a_\eta^*)$.

With $\hat{a}_{\eta+1}(a_\eta^*)$ as a new starting point, reiterate the above procedure for types $\eta+1$ and $\eta+2$ to pick $\hat{a}_{\eta+2}(a_\eta^*)$ in the intersection of the lower contour sets of types η and $\eta + 1$ with the upper contour sets of types $\eta + 1$ and $\eta + 2$: $\left[V_{\eta+1}(a_\eta^*) \cap L_\eta(a_\eta^*)\right] \cap$ $\left[V_{\eta+2}(\hat{a}_{\eta+1}(a_\eta^*)) \cap L_{\eta+1}(\hat{a}_{\eta+1}(a_\eta^*))\right]$. The intersection $V_{\eta+1}(a_\eta^*) \cap L_\eta(a_\eta^*) \cap V_{\eta+2}(\hat{a}_{\eta+1}(a_\eta^*))$ $\cap L_{\eta+1}(\hat{a}_{\eta+1}(a_\eta^*))$ is non-empty because $\hat{a}_{\eta+1}(a_\eta^*)$ can be picked arbitrarily close to a_η^* and type $\eta + 2$ is of higher rank than type η. Note that because $\hat{a}_{\eta+2}(a_\eta^*) \in L_{\eta+1}(\hat{a}_{\eta+1}(a_\eta^*))$ and $\hat{a}_{\eta+1}(a_\eta^*) \in L_\eta(\hat{a}_\eta(a_\eta^*))$, then $\hat{a}_{\eta+2}(a_\eta^*) \in L_\eta(\hat{a}_\eta(a_\eta^*))$, so η strictly prefers $\hat{a}_\eta(a_\eta^*)$ to $\hat{a}_{\eta+2}(a_\eta^*)$. Similarly, $\hat{a}_{\eta+2}(a_\eta^*) \in V_{\eta+2}(\hat{a}_\eta(a_\eta^*))$, so $\eta+2$ strictly prefers $\hat{a}_{\eta+2}(a_\eta^*)$ to $\hat{a}_\eta(a_\eta^*)$. Finally, $\eta+1$ strictly prefers $\hat{a}_{\eta+1}(a_\eta^*)$ to $\hat{a}_\eta(a_\eta^*)$ by the previous step in selecting $\hat{a}_{\eta+1}(a_\eta^*)$, and strictly prefers $\hat{a}_{\eta+1}(a_\eta^*)$ to $\hat{a}_{\eta+2}(a_\eta^*)$ by the current step in selecting $\hat{a}_{\eta+2}(a_\eta^*)$. Thus the sequence $\{\hat{a}_\eta(a_\eta^*), \hat{a}_{\eta+1}(a_\eta^*), \hat{a}_{\eta+2}(a_\eta^*)\}$ is incentive compatible for types $\eta, \eta + 1$ and $\eta + 2$.

With $\hat{a}_{\eta+2}(a_\eta^*)$ as a new starting point and by iterating the above procedure $\Theta - \eta - 2$ additional times, one picks a sequence of individual allocations $\{\hat{a}_\eta(a_\eta^*), \hat{a}_{\eta+1}(a_\eta^*), ..., \hat{a}_\Theta(a_\eta^*)\}$ which satisfy the definition of incentive compatibility in Definition 3 (see figure 3 for the case where there are two commodities in the economy). ∎

Corollary 2: Given an arbitrary type η and the first-best individual allocation a_η^*, there exists an incentive compatible set $\hat{A}(a_\eta^*)$ such that $u_\vartheta(a_\vartheta^*) > u_\vartheta(\hat{a}_\vartheta(a_\eta^*))$ for all $\vartheta > \eta$.

Proof: The result follows immediately from the fact that types are ranked so that $u_\vartheta(a_\vartheta^*) > u_\vartheta(a_\eta^*)$ whenever $\vartheta > \eta$ and the fact that, by following the algorithm of Lemma 2, the individual allocations in the set $\hat{A}(a_\eta^*)$ can be constructed arbitrarily close to a_η^*. ∎

Corollary 2 states that it is possible to design the out-of-equilibrium path individual allocations so that higher rank types strictly prefer their own first-best individual al-
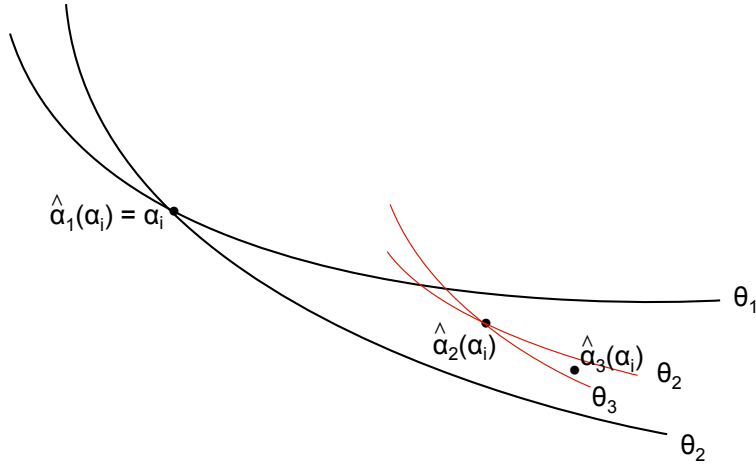
Figure 2: Incentive Compatible Individual Allocations for Two Goods and Three Types

locations to the former. The results in Lemmas 1, 2 and Corollaries 1,2 are useful in designing a sequential mechanism which implements the first-best allocation as a unique Perfect Bayesian Equilibrium. The mechanism consists of $\Theta - 1$ stages. Each stage is designed so as to incentivize the agents of a particular type to report truthfully and exit the mechanism, starting with the lowest rank type at stage one and rising to higher ranks in successive stages. At each stage the remaining agents are asked to report their type. If the number of reports for the type whose rank matches the number of the stage equals the interim number of agents with this type then these agents receive their respective first-best allocations and exit the mechanism while the rest of agents move to the next stage. Otherwise, they all receive incentive compatible individual allocations according to the type they reported and the mechanism ends. We show that designing off-the-equilibrium path individual allocations according to Lemma 2 and Corollary 2 is sufficient to induce higher rank agents to report truthfully irrespectively of the reports of lower ranks and this, in turn, forces the lower rank types to report truthfully as well. We thus obtain the result. The formal proof is provided below.

Theorem 1: If the realized distribution of types is common knowledge and preferences satisfy the single-crossing condition of Definition 2 then there exists a mechanism which implements any Pareto efficient allocation $\boldsymbol{a}^*$ which satisfies Anonymity as the

14

unique Perfect Bayesian Equilibrium. In this equilibrium agents report their private information truthfully.

Proof: Let $\boldsymbol{a}^*$ be a Pareto efficient allocation which satisfies Anonymity. Let a_ϑ^* be the individual first-best allocation, i.e. the individual allocation which the mechanism designer wishes to provide to an agent with type ϑ. Applying Lemma 2 on an individual allocation a_ϑ^* in $\boldsymbol{a}^*$, let $\hat{A}(a_\vartheta^*)$ be a set of *incentive compatible individual allocations* (according to notion incentive compatibility in definition 3 in page 12) for all types with rank higher or equal to ϑ with a_ϑ^* as the starting point (also recall that Lemma 1 allows one to rank types according to envy from the highest, Θ, to the lowest, 1). Moreover, let $\hat{A}(a_\vartheta^*)$ satisfy Corollary 2.

Given the above consider the following sequential mechanism with $\Theta - 1$ possible stages. Each stage is designed to aggregate the number of people who report a specific type (the 'stage-type'), starting with stage one, which is designed for type 1 and ascending up to stage $\Theta - 1$ which is designed for types $\Theta - 1$ and Θ. At each stage, the agents participating in it report their type. If the number of reports for the stage-type, $\mu(\vartheta)I$, matches the number of agents with this type, $\beta(\vartheta)I$, then agents who report it receive $a^*(\vartheta)$, and *exit* the mechanism. The remaining agents proceed to the next stage. If, however $\mu(\vartheta) \neq \beta(\vartheta)$, then all agents participating in the stage receive an individual allocation from the set $\hat{A}(a_\vartheta^*)$ according to the type they reported and the mechanism *ends*. In this case, agents who report a type with rank lower than ϑ receive a_ϑ^*. At the last stage, if $\mu(\Theta - 1) = \beta(\Theta - 1)$, then the agents who reported $\Theta - 1$ receive $a^*(\Theta - 1)$ and the remaining agents receive $a^*(\Theta)$ and the mechanism *ends*. If $\mu(\Theta - 1) \neq \beta(\Theta - 1)$, then agents receive individual allocations from the set $\hat{A}(a_{\Theta-1}^*)$, according to the type that they reported (agents who report a lower rank type than $\Theta - 1$ receive $a_{\Theta-1}^*$).

By backward induction, let us consider the last possible stage of the game, stage $\Theta - 1$. On any possible equilibrium path along the mechanism, stage $\Theta - 1$ is reached when only two types, $\Theta - 1$ and Θ, participate in it. To see that it is not possible to reach this stage with agents from other types participating, suppose the contrary, that is suppose that there is at least one agent of at least one type who participates in it. Because reaching stage $\Theta - 1$ requires that $\mu(\vartheta) = \beta(\vartheta)$ for all types with rank $\vartheta < \Theta - 1$, this means that there must be at least one agent of type Θ or $\Theta - 1$ who reported some other type κ in some previous stage (stage κ) of the mechanism, received individual allocation a_κ^* and *exited* the mechanism. However, such a strategy is not a

15

best-response for an agent of types Θ and $\Theta - 1$, because by reporting truthfully at stage κ he can improve his payoff. Specifically, if an agent of type Θ ($\Theta - 1$) reports truthfully in stage κ, instead of reporting κ, then $\mu(\kappa) = \beta(\kappa) - 1/I$ and he receives $\hat{a}_\Theta^*(a_\kappa^*)$ ($\hat{a}_{\Theta-1}^*(a_\kappa^*)$), which, by construction, he prefers to a_κ^*. Therefore, on equilibrium, only types $\Theta - 1$ and Θ participate in stage $\Theta - 1$. This reasoning also demonstrates that the number of agents participating in the last stage is exactly equal to the number of the two highest ranked types in the population, which is $[\beta(\Theta) + \beta(\Theta - 1)]I$. That is, all agents with type Θ and $\Theta - 1$ and only agents with these types participate in the last stage. Otherwise this stage would never be reached on the equilibrium path of the mechanism.

At stage $\Theta - 1$, reporting truthfully is a strictly dominant strategy for type Θ. To see this consider an agent of type Θ and suppose that all other agents report truthfully. Then, by reporting truthfully he receives a_Θ^*. If, instead, he misreports, then $\mu(\Theta-1) > \beta(\Theta - 1)$ and he receives $\hat{a}_{\Theta-1}(a_{\Theta-1}^*) \equiv a_{\Theta-1}^*$. Since Θ prefers a_Θ^* to $a_{\Theta-1}^*$, he prefers to report truthfully.

Another possibility is when a single agent of type $\Theta - 1$ misreports and the rest of agents report truthfully. If Θ reports truthfully in this case, then $\mu(\Theta - 1) < \beta(\Theta - 1)$ and Θ receives $\hat{a}_\Theta(a_{\Theta-1}^*)$. If, on the other hand, he misreports, then he covers for the misreporting of $\Theta - 1$ and $\mu(\Theta - 1) = \beta(\Theta - 1)$, in which case he receives $a_{\Theta-1}^*$. By construction, he prefers $\hat{a}_\Theta(a_{\Theta-1}^*)$ to $a_{\Theta-1}^*$ and therefore he prefers to report truthfully. The same reasoning applies to any case where there are multiple misreportings from both types, such that all but one cancel out and the report by a single Θ type determines whether $\mu(\Theta - 1) = \beta(\Theta - 1)$ (when he misreports) or whether $\mu(\Theta - 1) < \beta(\Theta - 1)$ (when he reports truthfully).

The final possible case to consider is when there are multiple misreports from (possibly) both types so that $\mu(\Theta - 1) < \beta(\Theta - 1) - 1/I$ so that a single Θ type's report can not cover the misreportings by the other agents. In this case agents receive individual allocations from the set $\hat{A}(a_{\Theta-1}^*)$. By construction, an agent of type Θ strictly prefers to report truthfully and receive $\hat{a}_\Theta(a_{\Theta-1}^*)$ instead of misreporting and receiving $\hat{a}_{\Theta-1}(a_{\Theta-1}^*)$. From the above three cases we conclude that reporting truthfully is a strictly dominant strategy for type Θ.

Conditional on this result, agents of type $\Theta - 1$ anticipate that all agents of type Θ report truthfully and their best response is to report truthfully as well. This is because if a single $\Theta - 1$ deviates then $\mu(\Theta - 1) < \beta(\Theta - 1)$ and he receives $\hat{a}_\Theta(a_{\Theta-1}^*)$, while if he reports truthfully he receives $a_{\Theta-1}^*$ which he prefers to $\hat{a}_\Theta(a_{\Theta-1}^*)$. And in the case where

16

there are multiple misreportings by other $\Theta - 1$ agents, still a $\Theta - 1$ agent prefers to report truthfully since, by construction, he prefers $\hat{a}_{\Theta-1}(a^*_{\Theta-1})$ to $\hat{a}_{\Theta}(a^*_{\Theta-1})$. Therefore, the unique Bayesian equilibrium of the last stage is for all agents to report truthfully.

Applying the same reasoning at stage $\Theta - 2$, the unique equilibrium of the stage is for agents of type $\Theta - 2$ to report truthfully while agents of types $\Theta - 1$ and Θ do not report type $\Theta - 2$ and continue to the final stage. To show this note that, as at stage $\Theta - 1$, the only types that participate at stage $\Theta - 2$ on the equilibrium path are $\Theta - 2$, $\Theta - 1$ and Θ and all agents with type $\Theta - 2$, $\Theta - 1$ or Θ participate in this stage.

Next, given the above analysis, the continuation value of the mechanism is $u_{\Theta}(a^*_{\Theta})$ for type Θ and $u_{\Theta-1}(a^*_{\Theta-1})$ for type $\Theta - 1$. Therefore, reporting type $\Theta - 2$ at stage $\Theta - 2$ is a strictly dominated strategy for these types. This is due to two arguments. The first is that if they believe that agents of type $\Theta - 2$ report truthfully, then it is a best-response for them to report either $\Theta - 1$ or Θ and proceed to the next stage than report $\Theta - 2$ and receive an individual allocation in the set $\hat{A}(a^*_{\Theta-2})$. The second effect is that if they believe that one or more agents of type $\Theta - 2$ misreport their type, then types Θ and $\Theta - 1$ prefer to report truthfully and receive the individual allocation which matches their type in the set $\hat{A}(a^*_{\Theta-2})$ than to misreport their type and receive any other individual allocation in the set (including $a^*_{\Theta-2}$). Therefore, no agent of type Θ or $\Theta - 1$ reports $\Theta - 2$ at stage $\Theta - 2$.

Given this, it is a best response for an agent of type $\Theta - 2$ to report truthfully irrespectively of the strategies of other agents with the same type. That is, suppose that all other $\Theta - 2$ agents report truthfully. Then an agent of type $\Theta - 2$ also prefers to report truthfully and receive $a^*_{\Theta-2}$ than to misreport his type and receive either $\hat{a}_{\Theta}(a^*_{\Theta-2})$ or $\hat{a}_{\Theta-1}(a^*_{\Theta-2})$. Also, if at least one other agent of type $\Theta - 2$ misreports his type, still it is a best-response for the other $\Theta - 2$ agents to report truthfully and receive $\hat{a}_{\Theta-2}(a^*_{\Theta-2}) \equiv a^*_{\Theta-2}$ than to misreport their type and receive either $\hat{a}_{\Theta}(a^*_{\Theta-2})$ or $\hat{a}_{\Theta-1}(a^*_{\Theta-2})$. Overall, given that types Θ and $\Theta - 1$ do not report $\Theta - 2$, it is best-response for $\Theta - 2$ to report truthfully. As a result, the unique equilibrium outcome of the stage is for type $\Theta - 2$ agents to receive $a^*_{\Theta-2}$ and exit the mechanism, while type Θ, $\Theta - 1$ agents proceed to the last stage.

The analysis of stages $\Theta - 2$ and $\Theta - 1$ proves that the continuation value of the mechanism for type $\Theta-2$ at a stage $\kappa < \Theta-2$ is $u_{\Theta-2}(a^*_{\Theta-2})$. Similarly, the continuation value for type $\Theta-1$ is $u_{\Theta-1}(a^*_{\Theta-1})$ and for type Θ is $u_{\Theta}(a^*_{\Theta})$. By induction, we conclude that the continuation value of the mechanism for a generic type ϑ at stage $\kappa < \vartheta$ is $u_{\vartheta}(a^*_{\vartheta})$. This means that the unique equilibrium outcome of an arbitrary stage κ of the

mechanism is for agents of type κ to report truthfully, receive individual allocation a_κ^* and exit the mechanism, while agents of types with higher rank do not report type κ and proceed to the next stage.

The first result required to prove this is that, on the equilibrium path, types of lower rank than κ do not participate in stage κ while all of the agents with higher or equal rank participate. Recall that stage κ is reached only if $\mu(\vartheta) = \beta(\vartheta)$ at all stages $\vartheta < \kappa$, which implies that $\sum_\vartheta \mu(\vartheta) = \sum_\vartheta \beta(\vartheta)$ for $\vartheta < \kappa$. If at least one agent of type $\eta < \kappa$ participates at stage κ (which means that he does not reveal his true type and he does not exit the mechanism at an earlier stage) then the requirement for stage κ to be reached $(\sum_\vartheta \mu(\vartheta) = \sum_\vartheta \beta(\vartheta))$ implies that there must be at least one agent of type $\lambda \geq \kappa$ who reports type η, receives a_η^* and exits the mechanism at stage η. But this strategy is not consistent with equilibrium-path play, for such an agent may switch his report from η to λ, induce $\mu(\eta) < \beta(\eta)$ and receive individual allocation $\hat{a}_\lambda(a_\eta^*)$, which, by construction, he prefers to a_η^*. Therefore, reaching stage κ on the equilibrium is only possible if no type with lower rank than κ participates in it.

This result in turn implies that any agent with type of lower rank than κ has exited the mechanism before stage κ and therefore all agents of types higher or equal to κ participate in the mechanism at stage κ. This is because the total number of agents who have exited the mechanism at stage κ is equal to $I(\sum_\vartheta \mu(\vartheta)) = I(\sum_\vartheta \beta(\vartheta))$ for $\vartheta < \kappa$, therefore the remaining number of agents who participate at stage κ is equal to $I(1 - \sum_\vartheta \beta(\vartheta))$ for $\vartheta < \kappa$, which is exactly equal to the number of agents with rank equal to or higher than κ. And since no agent of lower rank participates at stage κ then the number of participants in stage κ is equal to $I(1 - \sum_\vartheta \beta(\vartheta))$ only if all agents with rank equal to or higher than κ participate.

Next, it is a strictly dominant strategy for any type with $\lambda > \kappa$ not to report κ. First, suppose that all participants at stage κ report truthfully apart from agent i of type λ, who is considering the payoffs from his available strategies. If i reports κ then $\mu(\kappa) > \beta(\kappa)$, i receives $a_\kappa^* \equiv \hat{a}_\kappa(a_\kappa^*)$ and exits the mechanism. By not reporting κ he continues to the next stages of the mechanism, and, by applying the backward induction argument from the analysis of stages $\Theta-1$ and $\Theta-2$, he receives the individual allocation a_λ^*. By construction of these two individual allocations, it is a best-response for i not to report κ (recall that $u_\lambda(a_\lambda^*) > u_\lambda(\hat{a}_\lambda(a_\kappa^*)) > u_\lambda(a_\kappa^*)$ by Lemma 2 and Corollary 2).

Second, consider the possible case where one or more agents of type κ misreport their type. If i reports κ then he receives a_κ^* regardless of whether $\mu(\kappa) = \beta(\kappa)$ or not. But if he reports truthfully, then $\mu(\kappa) < \beta(\kappa)$ and he receives $\hat{a}_\lambda(a_\kappa^*)$, which is the

18

most preferable individual allocation in the set $\hat{A}(a_\kappa^*)$ for i by construction. Therefore, reporting κ is a strictly dominated strategy for any agent of type $\lambda > \kappa$. This implies that whether $\mu(\kappa) = \beta(\kappa)$ or not depends only on the reports of κ-type agents. This leads one to consider the available strategies for some agent, say j, of type κ. If all other agents of type κ report truthfully then j's report is pivotal in determining whether $\mu(\kappa) = \beta(\kappa)$ or not. If j reports truthfully, then $\mu(\kappa) = \beta(\kappa)$ and j receives a_κ^*. If j reports $\lambda \neq \kappa$ then $\mu(\kappa) < \beta(\kappa)$ and j receives $\hat{a}_\lambda(a_\kappa^*)$. By construction, j prefers to report truthfully is this case.

If, on the other hand, at least one other agent of type κ misreports his type, then $\mu(\kappa) < \beta(\kappa)$ irrespectively of j's report. In this case j's report determines which individual allocation in the set $\hat{A}(a_\kappa^*)$ he receives. Again, by construction, j prefers to report truthfully and receive individual allocation $a_\kappa^* \equiv \hat{a}_\kappa(a_\kappa^*)$. Therefore, regardless of the strategies of other κ-type agents, j prefers to report truthfully. As a result, the unique equilibrium outcome of the stage is agents of type κ to report truthfully, receive individual allocation a_κ^* and exit the mechanism, while agents of types with higher rank do not report type κ and proceed to the next stage. The result of theorem 1 follows by induction. ∎

Before concluding, we would like to briefly comment on the advantages our mechanism presents in comparison to the existing literature on implementation (see for example, Jackson, 1991, Maskin, 1999). First, our mechanism holds even with two agents (or even in the degenerate case of one agent). Second, the required message space is minimal, since agents send messages only about their own type. Third, we do not require any ad-hoc game, which has no equilibrium in pure strategies (like an integer game), in order to rule out undesirable equilibria. This is achieved by 'enticing' some of the misreporting agents to report truthfully, whenever there are multiple misrepresentations. Finally, even though the domain of preferences we consider is strictly smaller than in many other papers, still Assumptions 1 and 2 are relatively weak and there are many cases of interest that comply with them.

Conclusion

In this paper we consider a general hidden-type economy and, under relatively weak conditions, we show that it is possible to construct a mechanism which has a unique

Perfect Bayesian equilibrium, where all agents reveal their type truthfully and they receive a first-best individual allocation. If the realized frequencies of types are known, then one can aggregate the messages that all agents are sending out and uncover any misreport(s), even if the identity of the liar is not known.

Truth-telling, however, requires appropriately designed punishments for lying. If the punishment from detecting a lie is too severe, then some agents may deliberately lie about their type in order to force other agents to also do so. The lies cancel out and the former agents "steal" the allocations of the latter, who are forced to lie under the fear of extreme punishments. This can lead to coordination failures and multiplicity of equilibria. Therefore, uniqueness of the equilibrium requires a careful construction of the allocations when lies are detected. We show that such punishments exist when the single crossing condition holds.

References

Akerlof, G. (1970): "The Market for "Lemons": Quality Uncertainty and the Market Mechanism," The Quarterly Journal of Economics, 84, 488-500.

Athey, S., and I. Segal (2007): "Designing Efficient Mechanisms for Dynamic Bilateral Trading Games," American Economic Review, 97(2), 131-136.

Athey, S., and Segal, I. (2013): "An Efficient Dynamic Mechanism," Econometrica, 81(6), 2463-2485.

Armstrong, M. (1999): "Price Discrimination by a Many-Product Firm," Review of Economic Studies, 66:1, 151-168.

Ausubel, L. (2004): "An Efficient Ascending-Bid Auction for Multiple Objects," American Economic Review, 94, 1452-1475.

Ausubel, L. (2006): "An Efficient Dynamic Auction for Heterogeneous Commodities", American Economic Review, 96, 602-629.

Battaglini, M. (2005): "Long-Term Contracting with Markovian Consumers," The American Economic Review, 95(3), 637-658.

Bergemann, D., and J. Välimäki (2010): "The Dynamic Pivot Mechanism," Econometrica, 78(2), 771-789.

Casella, A. (2002): "Storable Votes," mimeo: Columbia University.

Clarke, E. (1971): "Multipart Pricing of Public Goods," Public Choice, 17-33.

Crémer, J., and R. McLean (1985): "Optimal Selling Strategies under Uncertainty for a Discriminating Monopolist when Demands are Interdependent," Econometrica, 53, 345-361.

Escobar, J. F., and Toikka, J. (2013): "Efficiency in Games with Markovian Private Information," Econometrica, 81(5), 1887-1934.

Gershkov, A., and B. Moldovanu (2009): "Learning about the Future and Dynamic Efficiency," The American Economic Review, 99(4), 1576-1587.

Groves, T. (1973): "Incentives in Teams," Econometrica, 617-631.

Jackson, M. (1991): "Bayesian Implementation," Econometrica, 59, 461-477.

Jackson, M. and H. Sonnnenschein (2007): "Overcoming Incentive Constraints by Linking Decisions," Econometrica, 75, 241-257.

Maskin, E. (1999): "Nash Equilibria and Welfare Optimality," Review of Economic Studies, 66, 23-38.

McAfee, R.P. (1992): "Amicable Divorce: Dissolving a Partnership with Simple Mechanisms," Journal of Economic Theory, 56:2, 266-293.

McLean, R., and A. Postlewaite (2002): "Informational Size and Incentive Compatibility," Econometrica, 70, 2421-2453.

McLean, R., and A. Postlewaite (2004): "Informational Size and Efficient Auctions," Review of Economic Studies, 71, 809-827.

Mezzetti, C. (2004): "Mechanism Design with Interdependent Valuations: Efficiency," Econometrica, 72, 1617-1626.

Pavan, A., Segal, I., and Toikka, J. (2014): "Dynamic Mechanism Design: A Myersonian Approach," Econometrica, 82(2), 601-653.

Perry, M. and P. Reny (2002): "An Efficient Auction," Econometrica, 70, 1199-1212.

Perry, M., and P. Reny (2005): "An Efficient Multi-Unit Ascending Auction," Review of Economic Studies, 72, 567-592.

Rothschild, M., and J. Stiglitz (1976): "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," Quarterly Journal of Economics, 90, 629-649.

Rustichini Aldo, Mark A. Satterthwaite, Steven R. Williams (1994), "Convergence to Efficiency in a Simple Market with Incomplete Information," Econometrica, 62, 1041-1063

Spence, M. (1973): "Job Market Signaling," The Quarterly Journal of Economics, 87, 355-374.

Vickrey, W. (1961): "Counterspeculation, Auctions and Competitive Sealed Tenders," Journal of Finance, p. 8-37.