UNIVERSITY *of* York

This is a repository copy of *Efficient allocations in economies with asymmetric information when the realized frequency of types is common knowledge*.

White Rose Research Online URL for this paper:
https://eprints.whiterose.ac.uk/id/eprint/149829/

Version: Published Version

Article:

White Rose
university consortium
Universities of Leeds, Sheffield & York

eprints@whiterose.ac.uk
https://eprints.whiterose.ac.uk/

CrossMark

Efficient allocations in economies with asymmetric information when the realized frequency of types is common knowledge

Aristotelis Boukouras[1] · Kostas Koufopoulos[2]

Abstract We consider a general economy, where agents have private information about their types. Types can be multidimensional 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.

1 Introduction

As first shown by Akerlof (1970), Spence (1973) and Rothschild and Stiglitz (1976), hidden-type (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 (see, e.g., Dréze et al. 2008). How are these restrictions relaxed as more information becomes common

✉ Aristotelis Boukouras
   Aris.Boukouras@le.ac.uk

   Kostas Koufopoulos
   kougop@unipi.gr

[1] University of Leicester, University Road, Leicester LE1 7RH, UK

[2] Department of Banking and Financial Management, Piraeus, Greece

🍃 Springer

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 the 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.

Furthermore, our mechanism can be useful in addressing real-world problems when agents' individual characteristics (types) are privately known, but the central authorities have good knowledge of the relevant aggregate numbers. For example, a government may know the aggregate output (income) and the productivity in each sector but not the productivity of each individual agent. In such a case, the government can employ our mechanism and taxes/subsidies to redistribute income without any losses in terms of efficiency. Another example is the identification of the amount of air pollution by each factory and the imposition of the corresponding penalty without direct inspection. For this result, our mechanism requires knowledge of the aggregate air pollution and the aggregate number of the machines used along with the machine specifications determining the emission of air pollutants.

The most closely related paper to ours is Piketty (1993) which also assumes that the realized distribution of types is common knowledge. Piketty considers the classical

taxation problem of Mirrlees (1971) and proposes a mechanism which allows the government to implement any first-best efficient allocation. Our model differs from Piketty in several respects: (i) Piketty considers only two private goods, while we allow for any finite number of goods (some of which could be public). (ii) We consider multidimensional asymmetric information instead of unidimensional. (iii) We allow for a more general feasibility set.

Jackson and Sonnnenschein (2007) 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.

Sun and Yannelis (2007a, 2008) show how incentive compatibility constraints can be overcome and first-best achieved in economies with private information and atomless agents. This efficiency result requires that the agents' private signals be conditionally independent (given the true state of nature) and utility functions and endowments be type independent. Also, this result does not hold for economies with asymmetric information and a fixed finite number of agents.

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 (Vickrey 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 quasi-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

---

[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.

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 et al. (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.

Following this direction of research, subsequent papers have examined the existence and efficiency of equilibria in economies with asymmetric information when: (i) there is a continuum of agents (Sun and Yannelis 2007b), (ii) there are infinitely many commodities (Podczeck and Yannelis 2008), (iii) there is an infinite number of agents and aggregate uncertainty (Qiao et al. 2016), (iv) there is a mix of "small" and "large" agents (Pesce 2010), and (v) there are more private goods than states of nature (Correia-da-Silva and Hervés-Beloso 2014). As in Rustichini et al. (1994), these papers restrict their focus to satisfying incentive compatibility constraints, while we are interested in full implementation.

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 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 multistage 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
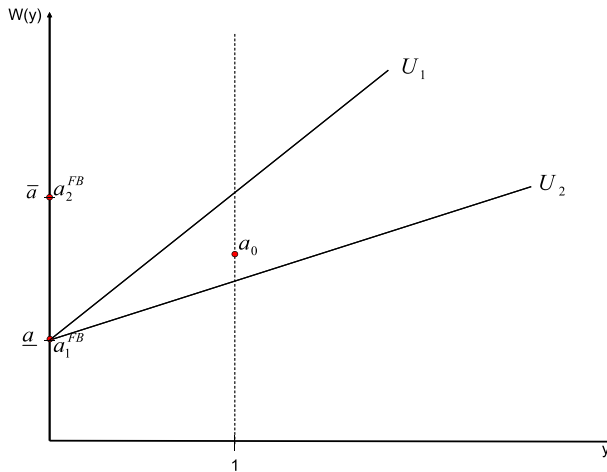
Fig. 1 Application to Spence (1973)

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 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.

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 Fig. 1) and those who report type 2 receive wage $w_2 = \overline{a}$ and zero education (contract α_2^{FB} in Fig. 1). In any other case, where the reported number of types does 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).

---

[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.

The above mechanism fully implements the first-best allocations in this economy. First, consider the strategies of type 2. It is clear that, irrespective of the reports of the other agents, it is a strictly dominant strategy for him to report his type truthfully. This 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 $\underline{a} + \epsilon - \frac{1}{\underline{a}}$ than to cover the lie by misreporting and receive $\underline{a}$, given that $\underline{a} + \epsilon - \frac{1}{\underline{a}} > \underline{a}$. 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 Fig. 1, contract a_0 denotes the offer to the workers, who report high productivity, when lies are detected.

3 The economy and definitions

The previous example was used 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 some useful definitions.

3.1 Agents, types and type profiles

The economy consists of a finite set $\mathcal{N}$ of agents. N stands for the aggregate number of agents. Θ is the finite set of potential types with ϑ denoting a generic element in the set. Θ denotes the aggregate number of types. In order to make our problem non-trivial we assume that $N \geq 2$ and $\Theta \geq 2$. $\omega \in \Theta^N$ is a *type profile*, i.e., a vector which contains a draw of a type for each of the N agents in the economy. ω is drawn from Θ^N by a generic probability distribution $f(\omega)$, that is $f : \Theta^N \to \Delta^{\Theta^N}$, where Δ^{Θ^N} is the unit simplex over Θ^N, i.e., $\{f \in \mathbf{R}_+ | \sum_{\omega \in \Theta^N} f(\omega) = 1\}$.

3.2 Information structure and realized frequency of types

It is assumed that each agent has private information about his own type, but does not know the types of the other agents, so that neither the agents nor the mechanism designer knows the realized type profile ω. However, both the agents and the mechanism designer know the *realized frequency of types* ϕ, where ϕ is a Θ-length vector with each element $\phi(\vartheta)$ denoting the proportion of agents who have type ϑ in the realized type profile ω. That is, ϕ denotes the relative frequency of each type, which materializes after types are drawn. Thus, $N(\vartheta) = \phi(\vartheta)N$ is the total number of agents with type ϑ. Since there are several type profiles which have the same realized frequency of types, let $\Omega(\phi)$ denote the set of all type profiles consistent with ϕ.

3.3 Feasible allocations

Suppose that there are L goods in total in the economy which are distributed to and consumed by the agents, with $L \geq 2$. $\mathcal{L}$ denotes the set of goods. At least two of the

goods in the economy are transferable and continuously divisible. If any of the goods is non-transferable, then it is a public good, i.e., all agents consume the same quantity of it.

An *individual allocation* a_i is a L-length vector that specifies a quantity for each good to be distributed to individual i. An *allocation* a is a $L \times N$-length vector that specifies an individual allocation for each agent.

Let $A(\mathcal{N}, \boldsymbol{\phi})$ be the set of all feasible allocations, with elements $a \in A(\boldsymbol{\phi}) \subseteq R_+^{L \times N}$. Note that we allow the feasibility set to depend on $\boldsymbol{\phi}$ but not on $\boldsymbol{\omega}$. That is, we implicitly assume that whenever two type profiles imply the same frequency of realized types, they also imply the same feasibility constraints. In other words, only the types of the agents play a role in determining which allocations are feasible, not their identities. For any subset of agents J, $J \subseteq \mathcal{N}$, we also define $A(J, \boldsymbol{\phi})$ to be the set of feasible allocations for the agents in J conditional on $\boldsymbol{\phi}$, that is $A(J, \boldsymbol{\phi}) \subseteq R_+^{L \times J}$.

Finally, note that later on in our analysis we focus on anonymous allocations that is on allocations where agents of the same type receive the same bundle of goods. Therefore, it is convenient to define individual allocations with respect to types rather than with respect to the agents' identities. To separate the two, we denote by a_ϑ the *type-specific individual allocation* of an agent with type ϑ, which is interpreted as the L-length vector of goods which the mechanism designer allocates to all agents of type ϑ.

3.4 Utility functions and contour sets

$u : R^L \times \Theta \to R$ is the *Bernoulli* utility, which is assumed to be continuous and differentiable with respect to all the L arguments that represent goods. Hence, $u_i(a_i; \vartheta)$ denotes the utility of agent i from consuming individual allocation a_i, conditional on i's type being ϑ. In the analysis that follows, however, we simplify the notation to $u_\vartheta(a_\vartheta)$ so that only an agent's type and her type-specific individual allocation show up.

Furthermore, we assume that the utility functions for all types satisfy the *single-crossing condition*. Formally, for any pair of types $\{\vartheta, \eta\} \in \Theta$, there exists at least one pair of goods $\{k, l\} \in \mathcal{L}$, such that $-\frac{\partial u_\vartheta / \partial a_{\vartheta l}}{\partial u_\vartheta / \partial a_{\vartheta k}} < -\frac{\partial u_\eta / \partial a_{\eta l}}{\partial u_\eta / \partial a_{\eta k}}$, where $a_{\vartheta k}$ is the individual allocation of good k to type ϑ. This condition is intuitive and is often used in the mechanism design literature.

The following definitions of lower and upper contour sets are also useful. $L_\vartheta(a_\vartheta)$ is the *lower contour set* of an agent with type ϑ associated with a type-specific individual allocation a_ϑ: $L_\vartheta(a_\vartheta) = \{c \in R_+^L : u_\vartheta(c) < u_\vartheta(a_\vartheta)\}$, and $V_\vartheta(a_\vartheta)$ is the *upper contour set* of type ϑ associated with a_ϑ: $V_\vartheta(a_\vartheta) = \{c \in R_+^L : u_\vartheta(c) > u_\vartheta(a_\vartheta)\}$.

3.5 Allocation rules, efficiency and anonymity

An *allocation rule* or *social choice rule* ρ is a mapping from the set of type profiles consistent with $\boldsymbol{\phi}$ to the set of feasible allocations, i.e., $\rho : \Omega(\boldsymbol{\phi}) \to A(\mathcal{N}, \boldsymbol{\phi})$.

An allocation a is *(Pareto) efficient* or *first-best* if there does not exist any other allocation $b \in A(\mathcal{N}, \boldsymbol{\phi})$ such that $u_i(b; \vartheta) > u_i(a; \vartheta)$ for at least one agent $i \in \mathcal{N}$

and $u_j(\boldsymbol{b}; \eta) \geq u_j(\boldsymbol{a}; \eta)$ for all other agents $j \in \mathcal{N} \backslash i$. Similarly, an allocation rule ρ is *(Pareto) efficient* or *first-best* if for every type profile $\boldsymbol{\omega} \in \Omega(\boldsymbol{\phi})$ it implements an efficient allocation $\boldsymbol{a}(\boldsymbol{\omega})$. In order to differentiate efficient allocations (allocation rules) from non-efficient ones we denote them by $\boldsymbol{a}^*(\rho^*)$. Thus, the definition of an efficient allocation rule can be formally expressed as $\rho^*(\boldsymbol{\omega}) = \boldsymbol{a}^*(\boldsymbol{\omega})$, $\forall$ $\boldsymbol{\omega} \in \Omega(\boldsymbol{\phi})$.

An allocation rule ρ is *anonymous* if, for every $\boldsymbol{\omega} \in \Omega(\boldsymbol{\phi})$, agents with the same type receive the same type-specific individual allocation. That is, every pair of agents $\{i, j\} \in \mathcal{N}$ that share the same type ϑ receives the same type-specific individual allocation $\boldsymbol{a}_\vartheta$. Thus, for allocation rules that satisfy *anonymity*, 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.

Note that the definition of the allocation rule implies that anonymity is a property of the final allocation as well. This means that if an allocation is anonymous, then all agents with the same type receive the same individual allocation. A direct consequence of this is that any anonymous allocation $\boldsymbol{a}$ can also be written as a Θ-length vector of type-specific individual allocations $\boldsymbol{a}_\vartheta$. Thus, if an allocation $\boldsymbol{a}^*$ is both efficient and anonymous, then it can be written as a Θ-length vector of *first-best type-specific individual allocations*, denoted by $\boldsymbol{a}_\vartheta^*$. Since our main results concern efficient and anonymous allocation rules, the use of this notation is very convenient in simplifying their presentation.

More importantly, if the allocations in a set $\{\boldsymbol{a}(\boldsymbol{\omega})\}$ are anonymous and they share the same realized distribution of types $\boldsymbol{\phi}$, then the set of type-specific individual allocations $\{\boldsymbol{a}_\vartheta\}$ for the respective types in Θ, as well the total number of bundles $\boldsymbol{a}_\vartheta$ to be distributed to the agents of each type, are fixed by $\boldsymbol{\phi}$ and they are independent of $\boldsymbol{\omega}$. This has important implications for the implementation of an efficient and anonymous allocation rule $\rho^*(\boldsymbol{\omega})$. From the perspective of individual agents, each final allocation $\boldsymbol{a}^*(\boldsymbol{\omega})$ is distinct from all the other because agents may receive different bundles to consume depending on their type. But, viewed from the perspective of types, all $\boldsymbol{a}^*(\boldsymbol{\omega})$ contain the same type-specific individual allocations and in the same numbers, they are just reshuffled across different agents. In other words, a mechanism designer who is tasked with implementing an anonymous $\rho^*(\boldsymbol{\omega})$ with a fixed $\boldsymbol{\phi}$ knows the number of agents of each type, and the bundle of goods each type should consume, but does not know who are the agents belonging to each type. Inducing agents to reveal this information truthfully is the main objective of implementation.

3.6 Mechanisms, incentive compatibility and self-selective sets

A *direct mechanism* $M_s(J_s, \boldsymbol{m}_s, h_s, g_s)$ is a game form with a sequence $s = \{1, 2, \ldots, S\}$ of stages, $S < +\infty$. $\mathcal{S}$ denotes the set of all stages. J_s is the set of

agents participating at stage s. $\boldsymbol{m}_s$ is a *message profile* , i.e., a report m_{is} of her type by every agent i in J_s. The set of all possible message profiles at stage s is $\mathcal{M}_s$, while $\boldsymbol{m}_{-i,s}$ denotes a type profile for all other agents apart from agent i. $h_s : \mathcal{M}_s \rightarrow J_{s+1}$, $J_{s+1} \subseteq J_s$, is a mapping from the set of message profiles to the set of agents who continue to stage $s + 1$. $g_s : \mathcal{M}_s \rightarrow A(J_s \setminus J_{s+1}, \boldsymbol{\phi})$ is a mapping from the set of message profiles to the set of allocations to the agents who exit the mechanism at stage s. If $J_{s+1} = J_s$, then $g_s = \emptyset$.

A sequence of message profiles $\{\boldsymbol{m}_s^* | s \in \mathcal{S}\}$ constitutes a *perfect Bayesian equilibrium* of the mechanism M_s if $\sum_{\boldsymbol{m}_{-i,s}^*} v_{is}(m_{is}^*, \boldsymbol{m}_{-i,s}^*)d_s(\boldsymbol{m}_{-i,s}^*) \geq \sum_{\boldsymbol{m}_{-i,s}^*} v_{is}(m_{is}, \boldsymbol{m}_{-i,s}^*)d_s(\boldsymbol{m}_{-i,s}^*)$, $\forall m_{is} \in \mathcal{M}_{is}, \forall i \in J_s, and \forall s \in \mathcal{S}$, where $v_{is}(\boldsymbol{m}_s)$ is the indirect utility of agent i at stage s under the message profile $\boldsymbol{m}_s$, $d_s(\boldsymbol{m}_{-i,s})$ is the probability that all other agents than i play the message profile $\boldsymbol{m}_{-i,s}$, and $\mathcal{M}_{is} \subseteq \Theta$ is agent's i message set at stage s.[3] Similarly, a mechanism M_s is *incentive compatible* if, for every agent i with type ϑ, $\sum_{\boldsymbol{\theta}_{-i}} v_{is}(\vartheta, \boldsymbol{\theta}_{-i})d_s(\boldsymbol{\theta}_{-i}) \geq \sum_{\boldsymbol{\theta}_{-i}} v_{is}(m_{is}, \boldsymbol{\theta}_{-i})d_s(\boldsymbol{\theta}_{-i}), \forall m_{is} \in \mathcal{M}_{is}, \forall i \in J_s, and \forall s \in \mathcal{S}$, where $\boldsymbol{\theta}_{-i}$ is the type profile of all agents apart from i. That is, a mechanism is incentive compatible when it is in the self-interest of every agent to report her type truthfully, conditional on the other agents reporting their type truthfully. Apart from the above common definitions, we also provide the following definitions of *self-selective type-specific individual allocation* and *self-selective set of individual allocations*, which are useful in proving our main result.

Given a set of types $J \subseteq \Theta$, the set $\hat{A}$ of type-specific individual allocations is a *self-selective set of type-specific individual allocations* or simply a *self-selective set*, if the set contains as many type-specific individual allocations as the number of types in the set J, and if, for each type $\vartheta \in J$, there exists a type-specific individual allocation $\hat{\boldsymbol{a}}_{\vartheta} \in \hat{A}$ such that $u_{\vartheta}(\hat{\boldsymbol{a}}_{\vartheta}) > u_{\vartheta}(\boldsymbol{a}_{\eta})$, $\forall \eta \in J \setminus \vartheta$. Any individual allocation that belongs to a *self-selective set* $\hat{A}$ is called a *self-selective type-specific individual allocation*. Simply put, the notion of self-selectiveness asserts that if a subset of types were left to choose a single type-specific individual allocation from a self-selective set, then each type would strictly prefer one of them over all other. This way agents truthfully reveal their type through their choice. This is useful for us because we use appropriately designed self-selective sets as out-of-equilibrium-path allocations in order to ensure the incentive compatibility of our mechanism.

3.7 Remarks

Overall, the economy is described by the following primitives: $E = \{\mathcal{N}, \Theta, A(\mathcal{N}, \boldsymbol{\phi}), \boldsymbol{\phi}, u\}$. This modeling of the economy is very flexible and can be easily extended to various dimensions. For example, the model allows for public goods problems, since some elements of the individual allocations can be common. Also, since we impose

---

[3] Therefore, for an agent i of type ϑ her indirect utility from stage s is $v_{is}(\boldsymbol{m}_s) = u_i(g_s(\boldsymbol{m}_s); \vartheta)$, if $i \notin J_{s+1}$, and $v_{is}(\boldsymbol{m}_s) = v_{i,s+1}(\boldsymbol{m}_{s+1})$, if $i \in J_{s+1}$, where the latter case states that i's payoff from stage s is simply her continuation payoff. Also note that, from i's perspective, the probability of the message profile $\boldsymbol{m}_{-i,s}$ is given by $d_s(\boldsymbol{m}_{-i,s}) = \prod_{j \in J_s \setminus i} \sum_{\eta \in \Theta} d_{js}(m_{js}, \eta) f_j(\eta | \boldsymbol{\phi}, \vartheta)$, where $d_{js}(m_{js}, \eta)$ is the probability that agent j plays m_{js} conditional on her type being η and $f_j(\eta | \boldsymbol{\phi}, \vartheta)$ is the probability that j is of type η conditional on $\boldsymbol{\phi}$ and on i's type.

no restrictions on $\boldsymbol{\phi}$ or the type-generating process that produces $\boldsymbol{\phi}$, types may or may not be independently distributed. Thus, several commonly used type-generating processes, such as Markov chains, are compatible with our formulation.

Furthermore, the model can be easily extended to include inter-dependent valuations for agents. For example, one may write the utility of agents as a function of both their own type and the type of other agents, $u_i(\boldsymbol{a}_i; \vartheta, \boldsymbol{\theta}_{-i})$. Or one may write the utility as a function of the realized distribution of types, $u_i(\boldsymbol{a}_i; \vartheta, \boldsymbol{\phi})$, if agents' identities do not play a role in other agents' utilities. In either case, as long as the effect of $\boldsymbol{\theta}_{-i}$ and $\boldsymbol{\phi}$ on u is identical across agents of the same type, so that the notation $u_\vartheta(.)$ correctly represents our model, all of the arguments and the results of the following section follow.

Finally, our model can be extended to economies with uncertainty. For example, let $\boldsymbol{\pi} : \boldsymbol{\phi} \rightarrow \Delta^Z$ be the probability distribution function over states, where Z the finite set of states and Δ^Z is the unit simplex $\{\boldsymbol{\pi} \in \boldsymbol{R}_+^Z | \sum_{z \in Z} \pi_z = 1\}$. In this case, $Y = Z \times L$, where L is the finite set of final goods and Y is the set of state-dependent goods. Thus, for a ϑ-type agent, her expected utility function is given by $u_\vartheta(\boldsymbol{a}_\vartheta; \boldsymbol{\phi}) = \sum_{z \in Z} x_\vartheta(\boldsymbol{a}_\vartheta, z)\pi_z(\boldsymbol{\phi})$, where $x_\vartheta(\boldsymbol{a}_\vartheta, z)$ is the decision-outcome payoff in state z. This formulation allows one to introduce uncertainty into the primitives of the economy while keeping the main framework unaffected, and thus our results apply equally well to it.

4 Implementation of first-best allocations

The main result of the paper is that any *anonymous efficient social choice rule* $\rho^*(\boldsymbol{\omega})$ can be uniquely implemented in perfect Bayesian equilibrium in a finite sequential mechanism, i.e., for each $\boldsymbol{\omega} \in \Omega(\boldsymbol{\phi})$, the unique perfect Bayesian equilibrium outcome of the mechanism is $\rho^*(\boldsymbol{\omega}) = \boldsymbol{a}^*(\boldsymbol{\omega})$.

In order to prove this result we proceed as follows. First, we show that Pareto efficiency implies a ranking of types according to envy (Lemma 1). This property is exploited in the design of the out-of-equilibrium-path allocations. Specifically, we show that the combination of the single-crossing condition with the result of Lemma 1 allows one to construct self-selective sets of type-specific individual allocations (Lemma 2). Each set contains type-specific individual allocations which are situated in the neighborhood of one of the first-best type-specific individual allocations $\boldsymbol{a}_\vartheta^*$ constituting $\boldsymbol{a}^*(\boldsymbol{\omega})$. The self-selective sets designed according to Lemma 2 are used as out-of-equilibrium-path allocations for each one of the stages of the mechanism. Finally, we combine these two results to prove our main theorem on the implementation of the efficient social choice rule. All proofs are provided in "Appendix."

We start the analysis with the result that any Pareto efficient and anonymous allocation $\boldsymbol{a}^*$ implies a ranking of types according to envy.

Lemma 1 *If $\boldsymbol{a}^*$ is a Pareto efficient and anonymous allocation, then, for any $\check{\Theta} \subseteq \Theta$ and the corresponding set of type-specific individual allocations $\check{A} = \{\boldsymbol{a}_\vartheta^* : \vartheta \in \check{\Theta}\}$, there exists at least one type $\vartheta \in \check{\Theta}$, who does not envy any of the other type-specific individual allocations in $\check{A}$: $u_\vartheta(\boldsymbol{a}_\vartheta^*) \geq u_\vartheta(\boldsymbol{a}_\eta^*), \ \forall \boldsymbol{a}_\eta^* \in \check{A}$.*

Lemma 1 allows one to construct a complete ranking of types according to envy. In particular, let $\mathcal{K} = \{\vartheta \in \Theta : u_\vartheta(a^*_\vartheta) \geq u_\eta(a^*_\eta), \ \forall \eta \in \Theta\}$ be the set of types who do not envy the type-specific individual allocation of any other type. By Lemma 1, we know that this set is non-empty. Then, by removing set $\mathcal{K}$ from the set of types Θ and applying Lemma 1 again to the remaining set, one obtains the set $\mathcal{K} - 1 = \{\vartheta \in \Theta : u_\vartheta(a^*_\vartheta) \geq u_\eta(a^*_\eta), \ \forall \eta \in \Theta - \mathcal{K}\}$. By iteration, one may define a total number of K envy sets, $\{\mathcal{K}, \mathcal{K} - 1, \mathcal{K} - 2, \ldots, 1\}$, $1 \leq K \leq \Theta$, ranked according to envy from the "highest" envy set $\mathcal{K}$ to the "lowest" set 1, such that the types in each one of them (say envy set $\mathcal{K} - k$) do not envy any of the types in their own set or lower sets (any $\mathcal{K} - l$, with $l \geq k$), but they envy some type(s) in a higher set (some type in an envy set $\mathcal{K} - l$, $l < k$).[4]

The K envy sets defined above could 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 envy sets. In order to simplify the required mechanism we rank types within each envy set that contains multiple types so that the mechanism is implemented over a full ranking of types. This can be achieved because the way these envy sets are constructed ensures that there is no envy between types which belong to the same envy set. For the rest of the analysis we assume that the mechanism designer ranks types according to the following simple rules:

1. Types who belong to a higher envy set are ranked above types who belong to a lower envy set.
2. If two types, ϑ and η, belong to the same envy set 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 set 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.

Two notes are worth making at this point. First, in order for the ranking of types to be complete, one should also examine the case where ϑ and η belong to the same envy set and $u_\vartheta(a^*_\vartheta) = u_\vartheta(a^*_\eta)$, $u_\eta(a^*_\eta) = u_\eta(a^*_\vartheta)$. However, this case is incompatible with single-crossing and so it is impossible to occur. Second, the ranking of types implied by rule 2 satisfies transitivity; hence, it is consistent.[5]

Overall, by following the above rules, the mechanism designer ranks all types according to envy from the type with the *lowest rank*, type 1, to the type with the *highest rank*, type Θ. From this point forward we use ϑ to denote the rank of a type, so that $1 \leq \vartheta \leq \Theta$. By the construction of this ranking, type Θ does not envy the first-best type-specific individual allocation of any other type. A generic type ϑ may envy the first-best type-specific individual allocation of a higher type ($\kappa > \vartheta$), but

---

[4] One extreme case is when an allocation exhibits no envy, in which case the set $\mathcal{K}$ contains the whole set of types and the total number of envy sets is only one (egalitarian allocations: $K = 1$). The other extreme case is when each envy set 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 and the total number of envy sets is equal to the total number of types ($K = \Theta$).

[5] The proof of this claim is provided in Appendix under Lemma 3. We are grateful to one anonymous referee for pointing it out to us.

does not envy the first-best type-specific individual allocation of any type with lower rank ($\eta < \vartheta$).[6]

The next step is to exploit the ranking of types and the single-crossing property to construct self-selective sets of type-specific individual allocations in the neighborhood of every first-best type-specific individual allocation contained in any $\boldsymbol{a}^*(\boldsymbol{\omega})$. These self-selective sets are used as out-of-equilibrium-path allocations in the mechanism, and they serve the purpose of "punishments" for misreporting types. To be more precise, given a complete ranking of types created by applying the rules above, for each type ϑ with respective first-best type-specific individual allocation $\boldsymbol{a}^*_\vartheta$, we construct a self-selective set of type-specific individual allocations $\hat{A}(\boldsymbol{a}^*_\vartheta)$ such that (i) it contains $\boldsymbol{a}^*_\vartheta$, and (ii) it contains a self-selective type-specific individual allocation for every type which has higher rank than type ϑ in an ϵ-neighborhood of $\boldsymbol{a}^*_\vartheta$, where ϵ is an arbitrarily small positive number. This means that any type with rank higher than ϑ prefers one of the type-specific individual allocations in $\hat{A}(\boldsymbol{a}^*_\vartheta)$ over all others. Lemma 2 below proves that the single-crossing condition allows one to construct self-selective sets with the above properties in the neighborhood of any $\boldsymbol{a}^*_\vartheta$.

Lemma 2 *Suppose that the single-crossing condition on utility functions holds and consider a first-best allocation $\boldsymbol{a}^*(\boldsymbol{\omega})$ with type-specific individual allocations $\boldsymbol{a}^*_\vartheta$, $\forall\, \vartheta \in \Theta$ and a complete ranking of types according to envy from the type with the lowest rank, type 1, to the type with the highest rank, type Θ. Then, for every $\boldsymbol{a}^*_\vartheta \in \boldsymbol{a}^*(\boldsymbol{\omega})$, $1 \le \vartheta \le \Theta - 1$, there exists $\epsilon > 0$ small enough and a set of self-selective type-specific individual allocations $\hat{A}(\boldsymbol{a}^*_\vartheta)$ such that:*

  (i) *$\boldsymbol{a}^*_\vartheta \in \hat{A}(\boldsymbol{a}^*_\vartheta)$*
  (ii) *there exists $\hat{\boldsymbol{a}}_\eta \in \hat{A}(\boldsymbol{a}^*_\vartheta)$ such that: $u_\eta(\hat{\boldsymbol{a}}_\eta) > u_\eta(\hat{\boldsymbol{a}}_\zeta)$, $\forall\, \hat{\boldsymbol{a}}_\zeta \in \hat{A}(\boldsymbol{a}^*_\vartheta)$,*
     $\hat{\boldsymbol{a}}_\zeta \ne \hat{\boldsymbol{a}}_\eta$, $\forall\, \eta : \vartheta \le \eta \le \Theta$
  (iii) *$\|\boldsymbol{a}^*_\vartheta - \hat{\boldsymbol{a}}_\eta\| \le \epsilon$, $\forall\, \hat{\boldsymbol{a}}_\eta \in \hat{A}(\boldsymbol{a}^*_\vartheta)$*
  (iv) *$u_\eta(\boldsymbol{a}^*_\eta) > u_\eta(\hat{\boldsymbol{a}}_\eta)$, $\forall\, \eta : \vartheta < \eta \le \Theta$*

Recall from Sect. 3 that the set of first-best type-specific individual allocations, $\boldsymbol{a}^*_\vartheta$, contained in $\boldsymbol{a}^*(\boldsymbol{\omega})$ is the same across all $\boldsymbol{\omega} \in \Omega(\boldsymbol{\phi})$. Therefore, any self-selective set with the properties of Lemma 2 is independent of $\boldsymbol{\omega}$, i.e., it retains its properties across all $\boldsymbol{\omega} \in \Omega(\boldsymbol{\phi})$.

The results in Lemmas 1 and 2 are useful in designing a sequential mechanism which implements the first-best allocation rule as a unique perfect Bayesian equilibrium. The proof of this result is provided in "Appendix." Here, we provide a verbal presentation of the mechanism and a sketch of the arguments used in the proof.

The mechanism consists of $\Theta - 1$ stages, one less than the total number of types, which is Θ. Each stage is designed so as to incentivize the agents of a particular type (the "target" type) to report truthfully their type, receive the respective first-best type-specific individual allocation for this type, and exit the mechanism. The first stage of the mechanism is designed for type 1, and each successive stage is designed for the

---

[6] Note that the ranking of types depends implicitly on the allocation rule ρ^*. But since the set of type-specific individual allocations $\{\boldsymbol{a}^*_\vartheta\}$ to be implemented by ρ^* does not depend on $\boldsymbol{\omega}$, the ranking of types is fixed for all $\boldsymbol{\omega} \in \Omega(\boldsymbol{\phi})$ as well.

type with rank one level higher than the type of the previous stage. The final stage, stage $\Theta - 1$, is designed for the two types with the highest rank, types $\Theta - 1$ and Θ.

In particular, at each stage, all agents who participate in it are asked to report their type. If the number of reports for the "target" type equals the realized number of agents with this type (which is known through ϕ), then the agents who report this type receive the respective first-best type-specific individual allocation of this type and exit the mechanism, while the agents, who report any other type, move to the next stage. If, on the other hand, the number of reports for the "target" type does not equal the realized number of agents with this type, then the mechanism terminates at this stage and all agents receive the self-selective type-specific individual allocation of the type they reported in that stage (i.e., their last report) from a set which has the properties of Lemma 2. Finally, if the mechanism does not terminate in any of the previous stages, the last stage of the mechanism is stage $\Theta - 1$.

The main reason that the above mechanism works is because, at each stage, types of higher rank than the "target" type strictly prefer to report truthfully their type than to report the "target" type, *irrespective* of the reports of the rest of the agents. In particular, if they expect the mechanism to end at the current stage, they prefer to report truthfully, and thus pick their most preferred type-specific individual allocation from the corresponding self-selective set triggered by the mechanism's termination, than to report any other type. And if they expect that the mechanism will not terminate at the current stage, then they prefer not to report the "target" type, because the continuation payoff from the mechanism is higher than exiting with the type-specific individual allocation of the "target" type. The last point holds because types of higher rank do not envy the "target" type's first-best type-specific individual allocation. Since higher rank types do not report this type, the best-response of agents of the "target" type is to report truthfully and by induction the result obtains.

Before the presentation of the theorem, a couple of remarks are noteworthy. First, the sequential nature of the mechanism used in the proof is not a necessary element of the result, i.e., one can find equivalent static mechanisms that implement the efficient allocation rule as a unique Bayesian equilibrium. Even though static mechanisms are less notation-heavy, they involve more case distinctions than our sequential mechanism, which makes the proof longer and more tedious. For the sake of elegance and transparency we opt for the sequential mechanism presented in detail in "Appendix."

Second, we remark again that, because $a^*(\omega)$ is anonymous and ϕ is common knowledge, the first-best type-specific individual allocations $\{a_{\vartheta}^*\}$ and the total number of bundles for each type are the same for every $a^*(\omega)$. It is only the identity of the agents who belong in each type that it is not known across each $\omega \in \Omega(\phi)$. Therefore, a mechanism implements $\rho^*(\omega)$ if it induces agents to report their type truthfully and if it allocates to them the respective first-best type-specific individual allocation of the type they report.

Theorem 1 *If the realized distribution of types is common knowledge and preferences satisfy the single-crossing condition, then there exists a mechanism which implements any Pareto efficient and anonymous allocation rule $\rho^*(\omega)$ as the unique perfect Bayesian equilibrium. In this equilibrium each agent reports her private information*

truthfully and receives the first-best type-specific individual allocation corresponding to her type.

Before concluding, we would like to briefly comment on the advantages our mechanism presents in comparison with the existing literature on implementation (see, e.g., 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 our assumptions on utility functions are relatively weak and there are many cases of interest that comply with them.

5 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 send 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 individual 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 individual allocations when lies are detected. We show that such punishments exist when the single-crossing condition holds.

Appendix

Proof of Lemma 1

First note that, by the assumptions made on the nature of the goods in this economy (i.e., either transferable private goods or public goods), it is possible to reallocate individual allocations from one agent to another. For example, it is possible for the mechanism designer to take a type-specific individual allocation $\boldsymbol{a}_\vartheta$ from agent i of

type ϑ and give it to agent j of type η, while j's type-specific allocation $\boldsymbol{a}_\eta$ may be reallocated to agent h and so on.

Then, consider a Pareto efficient and anonymous allocation $\boldsymbol{a}^*$ with type-specific individual allocations $\boldsymbol{a}_\vartheta^*$ and suppose that Lemma 1 does not hold for some subset $\check{\Theta}$ of types, $\check{\Theta} \subseteq \Theta$, with corresponding set of type-specific individual allocations $\check{A} = \{\boldsymbol{a}_\vartheta^* | \forall \vartheta \in \check{\Theta}\}$. This means that all types in $\check{\Theta}$ envy the type-specific individual allocation of at least one other type: $\forall \boldsymbol{a}_\vartheta^* \in \check{A}, \exists \eta \in \check{\Theta}, \eta \neq \vartheta : u_\vartheta(\boldsymbol{a}_\eta^*) > u_\vartheta(\boldsymbol{a}_\vartheta^*)$. But, since this holds for all types, then there exists at least one reassignment of the type-specific individual allocations among the agents with types in $\check{\Theta}$ such that some of them are made strictly better-off and the rest remain as well-off as under the original allocation $\boldsymbol{a}^*$.

For example, one could create a reassignment cycle by taking one agent from an arbitrary type, reassign to her one of the type-specific individual allocations that she envies, then proceed to the agent from whom the type-specific individual allocation was taken from, reassign to her one of the type-specific individual allocations that she envies, and so on, until the cycle returns back to one of the agents whose type-specific individual allocation has been reassigned. Because the set of agents is finite such a cycle is guaranteed to exist. Agents not involved in the cycle retain the type-specific individual allocations under $\boldsymbol{a}^*$.

However, the existence of such a cycle constitutes a Pareto improvement for a subset of agents in $\mathcal{N}$ and violates the initial assumption that $\boldsymbol{a}^*$ is Pareto efficient. $\qquad\square$

Proof of Lemma 2

The proof is done by the construction of a self-selective set that satisfies the properties of Lemma 2 with respect to an arbitrary type ϑ. For this purpose recall the definitions of upper contour set and lower contour set from Sect. 3, and consider a first-best and anonymous allocation $\boldsymbol{a}^*(\omega)$ and the corresponding first-best type-specific individual allocations $\boldsymbol{a}_\vartheta^*$, for any $\vartheta \in \Theta$. We use the single-crossing condition of utility functions and the complete ranking of types according to envy, which is possible as a result of Lemma 1 and which is assumed to satisfy the ranking rules 1–3 of page 10, to construct the required self-selective set. To show that this is possible, we take a type of an arbitrary rank ϑ, $1 \leq \vartheta \leq \Theta$ and consider the following algorithm for picking a sequence of type-specific individual allocations $\hat{\boldsymbol{a}}_\eta$ to be the elements of the self-selective set $\hat{A}(\boldsymbol{a}_\vartheta^*)$.

First, take $\boldsymbol{a}_\vartheta^*$ itself to be the first element of the set. Then use $\boldsymbol{a}_\vartheta^*$ as a starting point and pass the indifference curves of types ϑ and $\vartheta + 1$ from this point. By the single-crossing condition, the set $V_{\vartheta+1}(\boldsymbol{a}_\vartheta^*) \cap L_\vartheta(\boldsymbol{a}_\vartheta^*)$ is non-empty. For some $\epsilon > 0$ to be defined precisely later in the analysis, pick a type-specific individual allocation $\hat{\boldsymbol{a}}_{\vartheta+1}$ in the set $V_{\vartheta+1}(\boldsymbol{a}_\vartheta^*) \cap L_\vartheta(\boldsymbol{a}_\vartheta^*)$, such that $\|\boldsymbol{a}_\vartheta^* - \hat{\boldsymbol{a}}_{\vartheta+1}\| \leq \frac{\epsilon}{\Theta - \vartheta}$, to be the second element in $\hat{A}(\boldsymbol{a}_\vartheta^*)$. In the last expression, $\Theta - \vartheta$ is the natural number that represents the difference between the ranks of the type with the highest rank, Θ, with the rank of type ϑ. Also, because the utility functions are assumed to be continuous with respect to goods, it is always possible to find $\hat{\boldsymbol{a}}_{\vartheta+1}$ in the set $V_{\vartheta+1}(\boldsymbol{a}_\vartheta^*) \cap L_\vartheta(\boldsymbol{a}_\vartheta^*)$ such that

$\|\boldsymbol{a}^*_{\vartheta} - \hat{\boldsymbol{a}}_{\vartheta+1}\| \leq \frac{\epsilon}{\Theta-\vartheta}$ for any $\epsilon > 0$. Finally, note that, by construction, type ϑ strictly prefers $\boldsymbol{a}^*_{\vartheta}$ to $\hat{\boldsymbol{a}}_{\vartheta+1}$, while type $\vartheta + 1$ strictly prefers $\hat{\boldsymbol{a}}_{\vartheta+1}$ to $\boldsymbol{a}^*_{\vartheta}$.

With $\hat{\boldsymbol{a}}_{\vartheta+1}$ as a new starting point, reiterate the above procedure for types $\vartheta + 1$ and $\vartheta + 2$ to pick $\hat{\boldsymbol{a}}_{\vartheta+2}$ in the intersection of the lower contour sets of types ϑ and $\vartheta + 1$ with the upper contour sets of types $\vartheta + 1$ and $\vartheta + 2$, $\left[V_{\vartheta+1}(\boldsymbol{a}^*_{\vartheta}) \cap L_{\vartheta}(\boldsymbol{a}^*_{\vartheta})\right] \cap \left[V_{\vartheta+2}(\hat{\boldsymbol{a}}_{\vartheta+1}) \cap L_{\vartheta+1}(\hat{\boldsymbol{a}}_{\vartheta+1})\right]$, such that $\|\hat{\boldsymbol{a}}_{\vartheta+1} - \hat{\boldsymbol{a}}_{\vartheta+2}\| \leq \frac{\epsilon}{\Theta-\vartheta}$. Again, because of the continuity of the utility functions, the intersection $V_{\vartheta+1}(\boldsymbol{a}^*_{\vartheta}) \cap L_{\vartheta}(\boldsymbol{a}^*_{\vartheta}) \cap V_{\vartheta+2}(\hat{\boldsymbol{a}}_{\vartheta+1}) \cap L_{\vartheta+1}(\hat{\boldsymbol{a}}_{\vartheta+1})$ is non-empty and a type-specific individual allocation $\hat{\boldsymbol{a}}_{\vartheta+2}$ satisfying the above conditions exists. Note that, because $\hat{\boldsymbol{a}}_{\vartheta+2} \in L_{\vartheta+1}(\hat{\boldsymbol{a}}_{\vartheta+1})$ and $\hat{\boldsymbol{a}}_{\vartheta+1} \in L_{\vartheta}(\boldsymbol{a}^*_{\vartheta})$, then $\hat{\boldsymbol{a}}_{\vartheta+2} \in L_{\vartheta}(\boldsymbol{a}^*_{\vartheta})$, so ϑ strictly prefers $\boldsymbol{a}^*_{\vartheta}$ to $\hat{\boldsymbol{a}}_{\vartheta+2}$. Similarly, $\hat{\boldsymbol{a}}_{\vartheta+2} \in V_{\vartheta+2}(\hat{\boldsymbol{a}}_{\vartheta})$, so $\vartheta + 2$ strictly prefers $\hat{\boldsymbol{a}}_{\vartheta+2}$ to $\boldsymbol{a}^*_{\vartheta}$. Finally, $\vartheta + 1$ strictly prefers $\hat{\boldsymbol{a}}_{\vartheta+1}$ to $\boldsymbol{a}^*_{\vartheta}$ by the previous step in selecting $\hat{\boldsymbol{a}}_{\vartheta+1}$, and strictly prefers $\hat{\boldsymbol{a}}_{\vartheta+1}$ to $\hat{\boldsymbol{a}}_{\vartheta+2}$ by the current step in selecting $\hat{\boldsymbol{a}}_{\vartheta+2}$. Thus the sequence of type-specific individual allocations $\{\boldsymbol{a}^*_{\vartheta}, \hat{\boldsymbol{a}}_{\vartheta+1}, \hat{\boldsymbol{a}}_{\vartheta+2}\}$ constitutes a self-selective set for types ϑ, $\vartheta + 1$ and $\vartheta + 2$.

With $\hat{\boldsymbol{a}}_{\vartheta+2}$ as a new starting point and by iterating the above procedure $\Theta - \vartheta - 2$ additional times, one picks a sequence of type-specific individual allocations $\{\boldsymbol{a}^*_{\vartheta}, \hat{\boldsymbol{a}}_{\vartheta+1}, \ldots, \hat{\boldsymbol{a}}_{\Theta}\}$ which satisfies properties (i), (ii) and (iii) of Lemma 2. In particular, note that for any type η with rank $\vartheta < \eta \leq \Theta$, we have that $\|\boldsymbol{a}^*_{\vartheta} - \hat{\boldsymbol{a}}_{\eta}\| \leq \|\boldsymbol{a}^*_{\vartheta} - \hat{\boldsymbol{a}}_{\vartheta+1}\| + \sum_{k=1}^{\eta-\vartheta-1} \|\hat{\boldsymbol{a}}_{\vartheta+k} - \hat{\boldsymbol{a}}_{\vartheta+k+1}\| \leq (\eta - \vartheta)\frac{\epsilon}{\Theta-\vartheta} \leq \epsilon$, so property (iii) is verified (clearly, this property holds for $\boldsymbol{a}^*_{\vartheta}$ too).

It remains to define an appropriate ϵ such that property (iv) also holds. Define $\delta(\vartheta, \eta) = \min\{\|\boldsymbol{a}^*_{\vartheta} - \boldsymbol{a}_{\eta}\| \mid \forall \boldsymbol{a}_{\eta} \in I_{\eta}(\boldsymbol{a}^*_{\eta})\}$ to be the minimum distance between $\boldsymbol{a}^*_{\vartheta}$ and the indifference curve $I_{\eta}(\boldsymbol{a}^*_{\eta})$ of type η which passes through her first-best type-specific individual allocation. Define $\delta^{min}_{\vartheta} = \min\{\delta(\vartheta, \eta) \mid \forall \eta, \vartheta < \eta \leq \Theta\}$ to be the minimum distance between $\boldsymbol{a}^*_{\vartheta}$ and the indifference curves of all types with rank higher than ϑ which pass through their respective first-best type-specific individual allocations. Then for any $\epsilon < \delta^{min}_{\vartheta}$ property (iv) of Lemma 2 holds, because all the type-specific individual allocations in $\hat{A}(\boldsymbol{a}^*_{\vartheta})$ are at most distance ϵ away from $\boldsymbol{a}^*_{\vartheta}$. Therefore, they belong to $\bigcap_{k=1}^{\Theta-\vartheta} L_{\vartheta+k}(\boldsymbol{a}^*_{\vartheta+k})$, where $L_{\vartheta+k}(\boldsymbol{a}^*_{\vartheta+k})$ is the lower contour of type $\vartheta + k$, $k \geq 1$, that is associated with the first-best type-specific individual allocation $\boldsymbol{a}^*_{\vartheta+k}$.

Overall, the sequence of type-specific individual allocations $\{\boldsymbol{a}^*_{\vartheta}, \hat{\boldsymbol{a}}_{\vartheta+1}, \ldots, \hat{\boldsymbol{a}}_{\Theta}\}$ that constitutes $\hat{A}(\boldsymbol{a}^*_{\vartheta})$ is assumed to be constructed for some $\epsilon < \delta^{min}_{\vartheta}$ and hence properties (i)-(iv) of Lemma 2 hold. $\qquad\square$

The figure in the next page depicts the construction of a self-selective set that satisfies the properties of Lemma 2 for the case where there are two goods and three types, types 1, 2 and 3, with corresponding ranking, in the economy (Fig. 2).

Proof of transitivity of rule 2 on the complete ranking of types

Lemma 3 *Suppose that types ϑ, η, and ζ, with corresponding type-specific individual allocations $\boldsymbol{a}^*_{\vartheta}$, $\boldsymbol{a}^*_{\eta}$, and $\boldsymbol{a}^*_{\zeta}$, belong to the same envy set and suppose that: (i) $u_{\vartheta}(\boldsymbol{a}^*_{\vartheta}) >$*
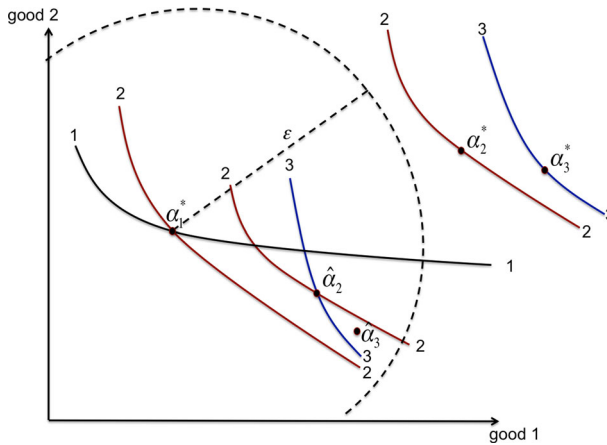
Fig. 2 A self-selective set for two goods and three types

$u_\vartheta(a_\eta^*), u_\eta(a_\vartheta^*) = u_\eta(a_\eta^*), (ii) u_\eta(a_\eta^*) > u_\eta(a_\zeta^*), u_\zeta(a_\eta^*) = u_\zeta(a_\zeta^*)$. *Then* $u_\vartheta(a_\vartheta^*) > u_\vartheta(a_\zeta^*)$.

Proof We prove the result by case elimination. There are two other cases possible, $u_\vartheta(a_\vartheta^*) < u_\vartheta(a_\zeta^*)$ and $u_\vartheta(a_\vartheta^*) = u_\vartheta(a_\zeta^*)$. We show that either of them is in contradiction with the assumption of single-crossing of the utility functions and the assumptions in the statement of Lemma 3.

(a) $u_\vartheta(a_\vartheta^*) < u_\vartheta(a_\zeta^*)$: This case is impossible to occur because it violates the assumption that all three types belong to the same envy set and therefore ϑ does not envy ζ.

(b) $u_\vartheta(a_\vartheta^*) = u_\vartheta(a_\zeta^*)$: This case contains two subcases, (i) $u_\zeta(a_\vartheta^*) = u_\zeta(a_\zeta^*)$ and (ii) $u_\zeta(a_\vartheta^*) < u_\zeta(a_\zeta^*)$. The subcase $u_\zeta(a_\vartheta^*) > u_\zeta(a_\zeta^*)$ is not examined for the same reason as the case $u_\vartheta(a_\vartheta^*) < u_\vartheta(a_\zeta^*)$.

(i) Suppose $u_\zeta(a_\vartheta^*) = u_\zeta(a_\zeta^*)$. But in this case, we have that both types ϑ and ζ are indifferent between their two first-best type-specific individual allocations a_ϑ^* and a_ζ^*, which violates the single-crossing assumption. Hence this case is impossible to hold.

(ii) Suppose $u_\zeta(a_\vartheta^*) < u_\zeta(a_\zeta^*)$. By $u_\eta(a_\vartheta^*) = u_\eta(a_\eta^*)$ and $u_\vartheta(a_\vartheta^*) > u_\vartheta(a_\eta^*)$, we have that $a_\eta^* \in I_\eta(a_\vartheta^*) \cap L_\vartheta(a_\vartheta^*)$, where $I_\eta(a_\vartheta^*)$ is the indifference curve of type η that passes though the type-specific individual allocation a_ϑ^* and $L_\vartheta(a_\vartheta^*)$ is the lower contour set of type ϑ associated with a_ϑ^*. Similarly, by $u_\vartheta(a_\vartheta^*) = u_\vartheta(a_\zeta^*)$ and $u_\eta(a_\eta^*) > u_\eta(a_\zeta^*)$, we have that $a_\zeta^* \in I_\vartheta(a_\vartheta^*) \cap L_\eta(a_\eta^*)$.

Therefore, by $u_\zeta(a_\eta^*) = u_\zeta(a_\zeta^*)$, we have that $I_\zeta(a_\zeta^*)$ passes through two points, a_η^* which belongs in the set $I_\eta(a_\vartheta^*) \cap L_\vartheta(a_\vartheta^*)$, and a_ζ^* which belongs to the set $I_\vartheta(a_\vartheta^*) \cap L_\eta(a_\eta^*)$. But any such indifference curve lies below a_ϑ^*, which means that $I_\zeta(a_\zeta^*) \in L_\zeta(a_\vartheta^*)$, which implies that $u_\zeta(a_\vartheta^*) > u_\zeta(a_\zeta^*)$ (see figure 3 for an illustration of the argument when $L = 2$). This violates the assumption that all three types belong to the same envy set and hence this case is also impossible to hold. Therefore, it must be that $u_\vartheta(a_\vartheta^*) > u_\vartheta(a_\zeta^*)$. □
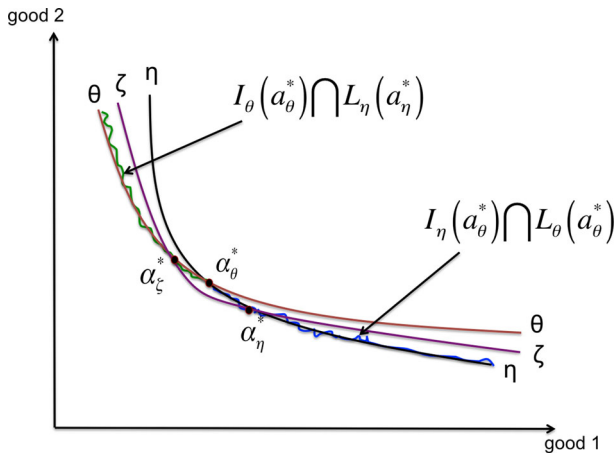
Fig. 3 The argument in the proof of Lemma 3 for the case of two goods

The result of Lemma 3 means that rule 2 in page 10, which is used in creating the ranking of types, is transitive. This is because, whenever the conditions of Lemma 3 apply to a triplet of types $\{\vartheta, \eta, \zeta\}$, then, as we proved $u_\vartheta(\boldsymbol{a}_\vartheta^*) > u_\vartheta(\boldsymbol{a}_\zeta^*)$, and $u_\vartheta(\boldsymbol{a}_\vartheta^*) > u_\vartheta(\boldsymbol{a}_\eta^*)$. As a result, rule 2 gives higher ranking to type ϑ over η and ζ and higher ranking to η over ζ. Thus, the ranking of types is consistent whenever the rules of page 10 are applied.

Proof of Theorem 1

Because the set of type-specific individual allocations $\{\boldsymbol{a}_\vartheta^*|\vartheta \in \boldsymbol{\Theta}\}$ and the corresponding number of bundles which constitute $\boldsymbol{a}^*(\boldsymbol{\omega})$ is independent of $\boldsymbol{\omega}$, the implementation of an efficient allocation rule $\boldsymbol{\rho}^*(\boldsymbol{\omega}) = \boldsymbol{a}^*(\boldsymbol{\omega})$ is equivalent to inducing agents to report their type truthfully and to receive the $\boldsymbol{a}_\vartheta^*$ that corresponds to their type. Thus, to prove Theorem 1, it is sufficient to find a mechanism that induces agents to reveal their type truthfully and which allocates to them the corresponding $\boldsymbol{a}_\vartheta^*$. The proof is done by construction. A sequential mechanism is presented and shown to generate this result. To facilitate its presentation, the proof is broken into several steps, which are summarized in Lemmas.

First, we reintroduce some notation used in the main text along with some new definitions, which are needed exclusively for the proof. Then, the mechanism is presented in formal terms. Subsequently, we provide a series of results on the best-responses of different types at different stages of the mechanism. Finally, these results are put together to provide the final result.

The following list contains the notation and definitions used in the proof of the results that follow:

– $\boldsymbol{a}^*(\boldsymbol{\omega})$ is the first-best and anonymous allocation to be implemented for every $\boldsymbol{\omega} \in \Omega(\boldsymbol{\phi})$

– The sequence of natural numbers $\{1, 2, \ldots, \vartheta, \ldots, \Theta\}$ represents the sequence of types, which are ranked according to envy and by applying the rules of page 10. Note that this sequence is invariant of $\boldsymbol{\omega}$

– After removing element Θ, the sequence $\{1, 2, \ldots, \vartheta, \ldots, \Theta - 1\}$ also represents the stages of the mechanism. Thus, stage ϑ refers to the ϑ-th stage of the mechanism, and the "target" type of the stage refers to type ϑ, the agents of which are expected to reveal their type truthfully at this stage

– The sequence $\{\boldsymbol{a}_1^*, \boldsymbol{a}_2^*, \ldots, \boldsymbol{a}_\vartheta^*, \ldots, \boldsymbol{a}_\Theta^*\}$ is the sequence of first-best type-specific individual allocations which corresponds to the sequence of types $\{1, 2, \ldots, \vartheta, \ldots, \Theta\}$. Note that this sequence is invariant of $\boldsymbol{\omega}$ too

– The sequence $\{\hat{A}(\boldsymbol{a}_1^*), \hat{A}(\boldsymbol{a}_2^*), \ldots, \hat{A}(\boldsymbol{a}_\vartheta^*), \ldots, \hat{A}(\boldsymbol{a}_\Theta^*)\}$ is the sequence of self-selective sets, which satisfy the properties of Lemma 2, and each one of them contains type-specific individual allocations in an ϵ-neighborhood of the respective first-best type-specific individual allocation

– $\hat{\boldsymbol{a}}_\eta(\boldsymbol{a}_\vartheta^*)$ denotes a generic element in the self-selective set $\hat{A}(\boldsymbol{a}_\vartheta^*)$, where η denotes the type who prefers this type-specific individual allocation over all others in the set, and $\boldsymbol{a}_\vartheta^*$ denotes the first-best type-specific individual allocation, in the ϵ-neighborhood of which $\hat{\boldsymbol{a}}_\eta(\boldsymbol{a}_\vartheta^*)$ is located

– Let $\overline{\boldsymbol{a}}_\eta(\vartheta) = \{\hat{\boldsymbol{a}}_\eta(\boldsymbol{a}_\zeta^*) \mid u_\eta(\hat{\boldsymbol{a}}_\eta(\boldsymbol{a}_\zeta^*)) \geq u_\eta(\hat{\boldsymbol{a}}_\eta(\boldsymbol{a}_\iota^*)), \ \forall \ \iota \in \Theta, \ \iota \leq \vartheta, \ \vartheta \leq \eta\}$ be the self-selective type-specific individual allocation which type η prefers over all other in the sequence of self-selective sets up to stage ϑ of the mechanism. If more than one type-specific individual allocations satisfy this definition, then let $\overline{\boldsymbol{a}}_\eta(\vartheta)$ be an arbitrary one among them

– Finally, recall the following definitions from Sect. 3.6. $m_{i\vartheta}$ denotes the "message" or "report" that agent i sends to the mechanism designer at stage ϑ of the mechanism. $\mu(\vartheta) = \{\#i \mid \mu_{i\vartheta} = \vartheta, i \in \mathcal{N}\}$ is the total number of agents who report type ϑ. $N(\vartheta)$ is the total number of agents in the economy with type ϑ, which is common knowledge since $N(\vartheta) = \phi(\vartheta)N$. J_ϑ is the set of agent who participate at stage ϑ. $g_{i\vartheta}(m_{i\vartheta}, \boldsymbol{m}_{-i,\vartheta})$ is the type-specific individual allocation that the mechanism specifies to agent i as a function of the message profile $\boldsymbol{m}_\vartheta = \{m_{i\vartheta}, \boldsymbol{m}_{-i,\vartheta}\}$ if i exits the mechanism at stage ϑ.

Lemma 4 *The following results are direct implications of the definition of* $\overline{a}_\eta(\vartheta)$

(i) $u_\eta(\overline{\boldsymbol{a}}_\eta(\vartheta)) > u_\eta(\hat{\boldsymbol{a}}_\zeta(\boldsymbol{a}_\iota^*)), \ \forall \, \hat{\boldsymbol{a}}_\zeta(\boldsymbol{a}_\iota^*) \in \hat{A}(\boldsymbol{a}_\iota^*), \ \forall \zeta \neq \eta, \ \forall \iota \leq \vartheta$
(ii) $\overline{\boldsymbol{a}}_\eta(\eta) = \boldsymbol{a}_\eta^*$
(iii) $u_\eta(\overline{\boldsymbol{a}}_\eta(\vartheta)) > u_\eta(\overline{\boldsymbol{a}}_\zeta(\vartheta)), \ \forall \{\zeta, \eta\} \in \Theta, \ \zeta \geq \vartheta, \ \eta \geq \vartheta, \ \zeta \neq \eta$
(iv) $u_\eta(\overline{\boldsymbol{a}}_\eta(\zeta)) \geq u_\eta(\overline{\boldsymbol{a}}_\eta(\vartheta)), \ \forall \zeta > \vartheta$

Proof (i) $u_\eta(\hat{\boldsymbol{a}}_\eta(\boldsymbol{a}_\iota^*)) > u_\eta(\hat{\boldsymbol{a}}_\zeta(\boldsymbol{a}_\iota^*)), \ \forall \zeta \neq \eta, \ \forall \iota \leq \vartheta$, by property (ii) of Lemma 2. Moreover, $u_\eta(\overline{\boldsymbol{a}}_\eta(\vartheta)) \geq u_\eta(\hat{\boldsymbol{a}}_\eta(\boldsymbol{a}_\iota^*)), \ \forall \iota \leq \vartheta$, by the definition of $\overline{\boldsymbol{a}}_\eta(\vartheta)$. These two inequalities together imply the result.

(ii) By the design of the self-selective set $\hat{A}(\boldsymbol{a}_\eta^*)$, $\boldsymbol{a}_\eta^*$ is the most preferable type-specific individual allocation for type η among all other elements in this set. In addition, for $\vartheta < \eta$, property (iv) of Lemma 2 and the definition of $\hat{\boldsymbol{a}}_\eta(\boldsymbol{a}_\vartheta^*)$ mean that $\boldsymbol{a}_\eta^*$ is strictly preferred to any other element of any self-selective set $\hat{A}(\boldsymbol{a}_\vartheta^*)$. These two facts together imply the result.

(iii) By construction, $\overline{a}_\zeta(\vartheta)$ is one of the type-specific individual allocations in the sequence $\{\hat{a}_\zeta(a_1^*), \hat{a}_\zeta(a_2^*), \ldots, \hat{a}_\zeta(a_\vartheta^*)\}$. Suppose it is the ι-th element in this sequence, with $1 \leq \iota \leq \vartheta$. Then, by property (ii) of Lemma 2 and the definition of $\overline{a}_\eta(\vartheta)$, we have that, $u_\eta(\overline{a}_\zeta(\vartheta)) = u_\eta(\hat{a}_\zeta(a_\iota^*)) < u_\eta(\hat{a}_\eta(a_\iota^*)) \leq u_\eta(\overline{a}_\eta(\vartheta))$.

(iv) $\overline{a}_\eta(\vartheta) = argmax\{u_\eta(a_\iota) \mid a_\iota \in \{\hat{a}_\eta(a_1^*), \hat{a}_\eta(a_2^*), \ldots, \hat{a}_\eta(a_\vartheta^*)\} \}$
$\overline{a}_\eta(\zeta) = argmax\{u_\eta(a_\iota) \mid a_\iota \in \{\hat{a}_\eta(a_1^*), \hat{a}_\eta(a_2^*), \ldots, \hat{a}_\eta(a_\zeta^*)\} \}$
$\{\hat{a}_\eta(a_1^*), \hat{a}_\eta(a_2^*), \ldots, \hat{a}_\eta(a_\vartheta^*)\} \subseteq \{\hat{a}_\eta(a_1^*), \hat{a}_\eta(a_2^*), \ldots, \hat{a}_\eta(a_\zeta^*)\}$

therefore $u_\eta(\overline{a}_\eta(\zeta)) \geq u_\eta(\overline{a}_\eta(\vartheta))$. □

The mechanism M^* We denote the mechanism by M^* to separate it from all other possible mechanisms. It consists of $\Theta - 1$ stages. In each stage ϑ, $1 \leq \vartheta \leq \Theta - 1$ each agent i who participates in the stage sends a report $m_{i\vartheta} \in \{\vartheta, \vartheta + 1, \ldots, \Theta\}$. That is, the strategy space of each agent is a report of a type at least as large in terms of ranking as the stage of the mechanism.

At an arbitrary stage ϑ, $1 \leq \vartheta \leq \Theta - 2$, if $\mu(\vartheta) = N(\vartheta)$, then agents with $m_{i\vartheta} = \vartheta$ receive the first-best type-specific individual allocation a_ϑ^* and *exit* the mechanism that is they consume a_ϑ^* immediately and they do not play in the remaining stages of the mechanism. The remaining agents, that is agents with $m_{i\vartheta} \neq \vartheta$, move to stage $\vartheta + 1$ of the mechanism. If $\mu(\vartheta) \neq N(\vartheta)$, then the mechanism *ends* at stage ϑ and each agent with $m_{i\vartheta} = \eta$, $\vartheta \leq \eta \leq \Theta$, receives and consumes the type-specific individual allocation $\overline{a}_\eta(\vartheta)$.

If the mechanism reaches stage $\Theta - 1$, the last stage possible, then, if $\mu(\Theta - 1) = N(\Theta - 1)$ agents with $m_{i,\Theta-1} = \Theta - 1$ receive and consume type-specific individual allocation $a_{\Theta-1}^*$, agents with $m_{i,\Theta-1} = \Theta$ receive and consume type-specific individual allocation a_Θ^*, and the mechanism *ends*. If $\mu(\Theta - 1) \neq N(\Theta - 1)$, then agents with $m_{i,\Theta-1} = \Theta - 1$ receive and consume type-specific individual allocation $a_{\Theta-1}^*$, agents with $m_{i,\Theta-1} = \Theta$ receive and consume type-specific individual allocation $\overline{a}_\Theta(\Theta - 1)$, and the mechanism *ends*. Formally, we have the following description of the mechanism M^*.

- For all ϑ such that $1 \leq \vartheta \leq \Theta - 2$:
  - If $\mu(\vartheta) = N(\vartheta)$, then for all $i \in \{m_{i\vartheta} = \vartheta\}$: $g_{i\vartheta}(\vartheta, \boldsymbol{m}_{-i,\vartheta}) = a_\vartheta^*$, $i \notin J_{\vartheta+1}$, and for all $i \in \{m_{i\vartheta} = \eta \neq \vartheta\}$: $g_{i\vartheta}(\eta, \boldsymbol{m}_{-i,\vartheta}) = \emptyset$, $i \in J_{\vartheta+1}$
  - If $\mu(\vartheta) \neq N(\vartheta)$, then for all $i \in \{m_{i\vartheta} = \eta\}$: $g_{i\vartheta}(\eta, \boldsymbol{m}_{-i,\vartheta}) = \overline{a}_\eta(\vartheta)$, for all $\eta \in \{\vartheta, \vartheta + 1, \ldots, \Theta\}$, and $J_{\vartheta+1} = \emptyset$,
- For $\vartheta = \Theta - 1$:
  - If $\mu(\Theta - 1) = N(\Theta - 1)$, then for all $i \in \{m_{i,\Theta-1} = \Theta - 1\}$: $g_{i,\Theta-1}(\Theta - 1, \boldsymbol{m}_{-i,\Theta-1}) = a_{\Theta-1}^*$, for all $i \in \{m_{i,\Theta-1} = \Theta\}$: $g_{i,\Theta-1}(\Theta, \boldsymbol{m}_{-i,\Theta-1}) = a_\Theta^*$, and $J_\Theta = \emptyset$
  - If $\mu(\Theta - 1) \neq N(\Theta - 1)$, then for all $i \in \{m_{i,\Theta-1} = \eta\}$: $g_{i,\Theta-1}(\eta, \boldsymbol{m}_{-i,\Theta-1}) = \overline{a}_\eta(\Theta - 1)$, for all $\eta \in \{\Theta - 1, \Theta\}$, and $J_\Theta = \emptyset$

Lemma 5 *Under M^*, at any stage ϑ, $1 \leq \vartheta \leq \Theta - 1$, it is a strictly dominated strategy for any agent of type η, with $\eta > \vartheta$, to report type ϑ.*

Proof Consider a generic stage ϑ, $1 \leq \vartheta \leq \Theta - 1$, and an agent i with a generic type η, $\vartheta < \eta \leq \Theta$. We show that the strategy $m_{i\vartheta} = \vartheta$ is strictly dominated by the

strategy $m_{i\vartheta} = \eta$, which is the truth-telling strategy, irrespective of i's beliefs whether the mechanism ends in stage ϑ or not.

In terms of how i's strategy affects the probability that the mechanism terminates at stage ϑ, there are three cases to consider. These three cases correspond to all the possible ways that the other agents' messages affect i's payoff.

(a) When all other agents reports, excluding i's report, are such $\mu(\vartheta) = N(\vartheta)$
(b) When all other agents reports, excluding i's report, are such $\mu(\vartheta) \geq N(\vartheta) + 1$
or $\mu(\vartheta) < N(\vartheta) - 1$
(c) When all other agents reports, excluding i's report, are such $\mu(\vartheta) = N(\vartheta) - 1$

(a) In this case i believes that the reports of the remaining agents are such that the number of agents who report type ϑ is exactly equal to the number of agents with type ϑ, given by the realized distribution of types $\boldsymbol{\phi}$. i is pivotal for the outcome of the mechanism in this case. If i reports type ϑ, then she increases $\mu(\vartheta)$ by one, so that, after including i's report, $\mu(\vartheta) = N(\vartheta)+1 \neq N(\vartheta)$, which induces the mechanism to end at stage ϑ. If, on the other hand, i reports truthfully her type η, then $\mu(\vartheta)$ remains unchanged after including her report. The mechanism proceeds to stage $\vartheta + 1$, and i along with it, since only agents with strategies $m_{j\vartheta} = \vartheta$, $j \neq i$, exit the mechanism at stage ϑ.

From the above two outcomes, i strictly prefers the second, which corresponds to the strategy $m_{i\vartheta} = \eta$, over the first, which corresponds to the strategy $m_{i\vartheta} = \vartheta$. To reach this conclusion, we first note that if i plays $m_{i\vartheta} = \vartheta$, then the mechanism terminates at stage ϑ and, by its rules, i receives and consumes $\bar{a}_\vartheta(\vartheta)$, which is equal to a_ϑ^* by Lemma 4(ii). If, on the other hand, i plays $m_{i\vartheta} = \eta$, her continuation payoff is at least as large $u_\eta(\bar{a}_\eta(\vartheta))$. The last argument is true because, by Lemma 4(iv), $u_\eta(\bar{a}_\eta(\zeta)) \geq u_\eta(\bar{a}_\eta(\vartheta))$, $\forall \zeta > \vartheta$. Finally, because $u_\eta(\bar{a}_\eta(\vartheta)) > u_\eta(\bar{a}_\vartheta(\vartheta)) = u_\eta(a_\vartheta^*)$ by Lemma 4(iii), reporting truthfully and proceeding to stage $\vartheta + 1$ is strictly preferable to i than terminating the mechanism at stage ϑ.

(b) In the second case i believes that the reports of the remaining agents are such that the number of agents who report type ϑ does not equal the number of agents with type ϑ, irrespective of her own report. That is, i expects that $\mu(\vartheta) \neq N(\vartheta)$ for all her strategies $m_{i\vartheta}$, and, therefore, she believes that the mechanism terminates at stage ϑ and she cannot affect this outcome. Under this belief, i expects to receive and consume $\bar{a}_\eta(\vartheta)$ if she reports η and she expects to receive and consume $\bar{a}_\vartheta(\vartheta)$ if she reports ϑ. By Lemma 4, $u_\eta(\bar{a}_\eta(\vartheta)) > u_\eta(\bar{a}_\vartheta(\vartheta)) = u_\eta(a_\vartheta^*)$, and therefore, i strictly prefers to report truthfully than to report type ϑ.

(c) In the last case i believes that the reports of the remaining agents are such that the number of agents who report type ϑ does not equal the number of agents with type ϑ by only one. If i reports ϑ, then, including her report, $\mu(\vartheta) = N(\vartheta)$ and so the mechanism does not terminate at stage ϑ. If i reports η, then $\mu(\vartheta) \neq N(\vartheta)$ and the mechanism terminates at stage ϑ. Therefore i's report is pivotal in determining the outcome of the mechanism for the other agents. For i, however, the comparison of payoffs is the same as in case (b). If i reports ϑ, then she receives and consumes a_ϑ^* and exits the mechanism. If, on the other hand, she reports η the mechanism terminates at stage ϑ and she receives and consumes $\bar{a}_\eta(\vartheta)$. As in case (b), $u_\eta(\bar{a}_\eta(\vartheta)) > u_\eta(a_\vartheta^*)$, by Lemma 4, and therefore i strictly prefers to report truthfully.

Overall, under all possible beliefs regarding the outcomes of the mechanism, i strictly prefers reporting her type truthfully over reporting type ϑ, and, therefore, the last strategy is strictly dominated for her and will never be played in any equilibrium of the mechanism. □

Lemma 6 *If stage ϑ of M^* is reached, $\vartheta > 1$, then all agents with types of rank lower than type ϑ have exited the mechanism at previous stages, and all agents with types of rank equal or higher than ϑ participate in stage ϑ.*

Proof The proof is done by induction. Consider stage 1 of M^*. By the design of M^*, moving from stage 1 to stage 2 happens only if $\mu(1) = N(1)$. Since, by Lemma 5, all agent with types $\{2, 3, \ldots \Theta\}$ do not report type 1 at stage 1, the requirement for proceeding to stage 2 is met only if all agents with type 1 report their type truthfully and exit the mechanism at stage 1. Therefore, if stage 2 is reached, all the agents of type 2 or higher participate, but no agent of type 1 participates.

Similarly, stage 3 is reached only if $\mu(2) = N(2)$ at stage 2. By Lemma 5, no agent of type 3 or of higher type reports type 2 so stage 3 is reached only if all agents of type 2 report their type truthfully. Thus, if stage 3 is reached, then all agents of type 2 have exited at stage 2 and all agents of type 3 or higher participate in stage 3. By applying the same reasoning to each successive stage, the result obtains. □

Lemma 7 *Under M^*, the unique best-response of any agent with type ϑ when she participates in stage ϑ is to report her type truthfully.*

Proof By Lemma 6, in stage ϑ participate all agents of type ϑ or higher and no agent of rank lower than ϑ. By Lemma 5, all agents of type $\vartheta + 1$ or higher do not report type ϑ. Therefore, for an agent i of type ϑ there are two distinct sets of beliefs consistent with the equilibrium behavior of other types.

  (a) When all other agents of type ϑ report truthfully
  (b) When one or more agents of type ϑ do not report truthfully

(a) In this case, excluding i's report, $\mu(\vartheta) = N(\vartheta) - 1$. If i reports $m_{i\vartheta} = \vartheta$, then, after including her report, $\mu(\vartheta)$ increases by one so that $\mu(\vartheta) = N(\vartheta)$, the mechanism proceeds to stage $\vartheta + 1$ if $\vartheta < \Theta - 1$ or terminates if $\vartheta = \Theta - 1$, and i receives and consumes $\boldsymbol{a}_\vartheta^*$. If i reports $m_{i\vartheta} = \eta \neq \vartheta$, then, after including her report, $\mu(\vartheta) < N(\vartheta)$, the mechanism terminates at stage ϑ, and i receives and consumes $\overline{\boldsymbol{a}}_\eta(\vartheta)$. By Lemma 4, $u_\vartheta(\boldsymbol{a}_\vartheta^*) = u_\vartheta(\overline{\boldsymbol{a}}_\vartheta(\vartheta)) > u_\vartheta(\overline{\boldsymbol{a}}_\eta(\vartheta))$, therefore, i strictly prefers to report truthfully than to report any other type.

(b) In this case, excluding i's report, $\mu(\vartheta) \leq N(\vartheta) - 2$, and, therefore, irrespective of i's report, the mechanism terminates at stage ϑ. If i reports $m_{i\vartheta} = \vartheta$ she receives and consumes $\overline{\boldsymbol{a}}_\vartheta(\vartheta) = \boldsymbol{a}_\vartheta^*$ and if she reports $m_{i\vartheta} = \eta \neq \vartheta$, then she receives and consumes $\overline{\boldsymbol{a}}_\eta(\vartheta)$. As with case (a) above, $u_\vartheta(\overline{\boldsymbol{a}}_\vartheta(\vartheta)) > u_\vartheta(\overline{\boldsymbol{a}}_\eta(\vartheta))$, and, therefore, i strictly prefers to report truthfully than to report any other type.

Overall, under any set of beliefs regarding the possible strategies of the other agents with type ϑ, i strictly prefers to report truthfully than report any other type and the result of the Lemma obtains. □

Together, the construction of M^* along with the results of Lemmas 5, 6 and 7, implies the result of Theorem 1 in page 12. To be precise, at any stage ϑ of M^*, $1 \leq \vartheta \leq \Theta - 1$, agents of type with rank higher than ϑ do not report type ϑ, by Lemma 5, agents of type with rank lower than ϑ have exited the mechanism at earlier stages, by Lemma 6, and the unique best-response of agents with type ϑ is to report their type truthfully, receive a_ϑ^* and exit the mechanism, by Lemma 7.

The result of Lemma 7 also implies that, since $\mu(\vartheta) = N(\vartheta)$, $1 \leq \vartheta \leq \Theta - 1$, on the equilibrium path, all stages of the mechanism are reached. Moreover, since Lemmas 5 and 7 apply to the last stage as well, stage $\Theta - 1$, then $\mu(\Theta - 1) = N(\Theta - 1)$. Therefore, by the design of M^*, all agents of type $\Theta - 1$ receive and consume $a_{\Theta-1}^*$, all agents of type Θ receive and consume a_Θ^* and the mechanism terminates. Thus, the unique perfect Bayesian equilibrium outcome of M^* is for all agents to report their type truthfully and to receive the first-best type-specific individual allocation corresponding to their type, as Theorem 1 requires. □

References

Akerlof, G.: The market for "Lemons": quality uncertainty and the market mechanism. Q. J. Econ. **84**, 488–500 (1970)

Athey, S., Segal, I.: Designing efficient mechanisms for dynamic bilateral trading games. Am. Econ. Rev. **97**, 131–136 (2007)

Athey, S., Segal, I.: An efficient dynamic mechanism. Econometrica **81**, 2463–2485 (2013)

Armstrong, M.: Price discrimination by a many-product firm. Rev. Econ. Stud. **66**, 151–168 (1999)

Ausubel, L.: An efficient ascending-bid auction for multiple objects. Am. Econ. Rev. **94**, 1452–1475 (2004)

Ausubel, L.: An efficient dynamic auction for heterogeneous commodities. Am. Econ. Rev. **96**, 602–629 (2006)

Battaglini, M.: Long-term contracting with Markovian consumers. Am. Econ. Rev. **95**, 637–658 (2005)

Bergemann, D., Välimäki, J.: The dynamic pivot mechanism. Econometrica **78**, 771–789 (2010)

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

Clarke, E.: Multipart pricing of public goods. Public Choice **11**, 17–33 (1971)

Correia-da-Silva, J., Hervés-Beloso, C.: Irrelevance of private information in two-period economies with more goods than states of nature. Econ. Theory **55**, 439–455 (2014)

Crémer, J., McLean, R.: Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent. Econometrica **53**, 345–361 (1985)

Dréze, J., Minelli, E., Tirelli, M.: Production and financial policies under asymmetric information. Econ. Theory **35**, 217–231 (2008)

Escobar, J.F., Toikka, J.: Efficiency in games with Markovian private information. Econometrica **81**, 1887–1934 (2013)

Gershkov, A., Moldovanu, B.: Learning about the future and dynamic efficiency. Am. Econ. Rev. **99**, 1576–1587 (2009)

Groves, T.: Incentives in teams. Econometrica **41**, 617–631 (1973)

Jackson, M.: Bayesian implementation. Econometrica **59**, 461–477 (1991)

Jackson, M., Sonnnenschein, H.: Overcoming incentive constraints by linking decisions. Econometrica **75**, 241–257 (2007)

Maskin, E.: Nash equilibria and welfare optimality. Rev. Econ. Stud. **66**, 23–38 (1999)

McAfee, R.P.: Amicable divorce: dissolving a partnership with simple mechanisms. J. Econ. Theory **56**, 266–293 (1992)

McLean, R., Postlewaite, A.: Informational size and incentive compatibility. Econometrica **70**, 2421–2453 (2002)

McLean, R., Postlewaite, A.: Informational size and efficient auctions. Rev. Econ. Stud. **71**, 809–827 (2004)

Mezzetti, C.: Mechanism design with interdependent valuations: efficiency. Econometrica **72**, 1617–1626 (2004)

Mirrlees, J.: An exploration in the theory of optimum income taxation. Rev. Econ. Stud. **38**, 175–208 (1971)

Pavan, A., Segal, I., Toikka, J.: Dynamic mechanism design: a Myersonian approach. Econometrica **82**, 601–653 (2014)

Perry, M., Reny, P.: An efficient auction. Econometrica **70**, 1199–1212 (2002)

Perry, M., Reny, P.: An efficient multi-unit ascending auction. Rev. Econ. Stud. **72**, 567–592 (2005)

Pesce, M.: On mixed markets with asymmetric information. Econ. Theory **45**, 23–53 (2010)

Piketty, T.: Implementation of first-best allocations via generalized tax schedules. J. Econ. Theory **61**, 23–41 (1993)

Podczeck, K., Yannelis, N.C.: Equilibrium theory with asymmetric information and with infinitely many commodities. J. Econ. Theory **141**, 152–183 (2008)

Qiao, L., Sun, Y., Zhang, Z.: Conditional exact law of large numbers and asymmetric information economies with aggregate uncertainty. Econ. Theory **62**, 43–64 (2016)

Rothschild, M., Stiglitz, J.: Equilibrium in competitive insurance markets: an essay on the economics of imperfect information. Q. J. Econ. **90**, 629–649 (1976)

Rustichini, A., Satterthwaite, A.M., Williams, R.S.: Convergence to efficiency in a simple market with incomplete information. Econometrica **62**, 1041–1063 (1994)

Spence, M.: Job market signaling. Q. J. Econ. **87**, 355–374 (1973)

Sun, Y.N., Yannelis, N.C.: Perfect competition in asymmetric information economies: compatibility of efficiency and incentives. J. Econ. Theory **134**, 175–194 (2007a)

Sun, Y.N., Yannelis, N.C.: Equilibria and incentives in large asymmetric information economies. Games Econ. Behav. **61**, 131–155 (2007b)

Sun, Y.N., Yannelis, N.C.: Ex ante efficiency implies incentive compatibility. Econ. Theory **36**, 35–55 (2008)

Vickrey, W.: Counterspeculation, auctions and competitive sealed tenders. J. Finance **17**, 8–37 (1961)