Information Aggregation and Adverse Selection*

Aristotelis Boukouras† Kostas Koufopoulos‡

This Version: 03 June 2011
First Version: 24 July 2007

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 relative frequency of types in the population (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 proof is done by iterated elimination of strictly dominated strategies and hence our equilibrium is also Bayes-Nash equilibrium. This result requires weak restrictions on preferences (Local Non-Common Indifference Property) and on the Pareto correspondence (Anonymity) and it is robust to small perturbations regarding the frequency of types. Our mechanism is particularly useful for games with large populations and independently distributed types, where the realized frequency of types is the ex-ante distribution. In these cases, full implementation of efficient allocations does not require any additional information than what is usually assumed.

Keywords: adverse selection, anonymity, first-best allocations, full implementation, information aggregation, mechanism design, single-crossing property, Pareto correspondence
JEL Classification: D71, D82, D86

---

*We would like to thank Luis Braido, Theodoros Diasakos, Peter Hammond, Christopher Hennessy, Matthew Jackson, Eric Maskin, Claudio Mezzetti, Chris Mürris, Asako Ohinata, Motty Perry, Philip Reny, Paul Youdell and Nikos Zygouras, for their very useful comments and remarks, for which we feel grateful.

†Courant Research Centre Poverty, Equity and Growth, Georg-August University Göttingen, aboukou@gwdg.de

‡University of Warwick, Warwick Business School, Kostas.Koufopoulos@wbs.ac.uk

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[1]. 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 towards 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 relative frequency of types), then it is possible, under fairly general conditions, to implement first-best allocations. More precisely, we consider an economy with asymmetric information, where each agent has private information about his type. We also assume that i) the relative frequency of types is common knowledge, ii) preferences satisfy the Local Non-Common Indifference Property and iii) the social choice rule satisfies Anonymity. Given these general 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. We obtain our equilibrium by using iterated elimination of strictly dominated strategies, and hence it is also a Bayes-Nash equilibrium.

This result has two interpretations. On one hand, 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 information 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 can he do? Can he induce the customers to reveal the orders they have made truthfully 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.

On the other hand, one can interpret this result as an application of the law of large numbers. If the ex-ante probability distribution is known, then, for sufficiently large populations, one can obtain a quite accurate estimate of the aggregate number of

[1]The title of our paper may be slightly misleading. Adverse selection is, of course, the outcome that may be generated in private information environments. The true source of the problem is the hidden information. Despite the fact that in our paper we have a hidden-types economy, we show that in the equilibrium of our mechanism, individuals reveal their information truthfully and they receive first-best allocations based on that. Therefore, adverse selection problems never arise as an equilibrium of our game. So, our main claim is that information aggregation, under certain conditions, can eliminate the possibility of adverse selection outcomes.

agents who have a specific type and, based on this information, he can address adverse selection problems. An example of this case would be insurance companies, which have data on millions of cases, collected over decades, and know with very high accuracy the probability of certain accidents taking place and how personal characteristics affect these probabilities. While the main result is originally stated for the case where the frequency of types is known with perfect precision, we subsequently prove that it also holds in the case where it is known with a small noise.

Our formulation is general enough to accommodate both interpretations and the intuition behind the result is common. If the 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. As a consequence, 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 extreme. 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 in terms of the aggregate information 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 allocations when lies are detected. We show that such punishments exist when the indifference curves of different types are not locally identical, meaning that in the neighborhood of any allocation one can find other allocations such that each type prefers one of these over the rest.

We acknowledge that requiring the realized frequency of each type being common knowledge seems a very restrictive assumption, particularly if it concerns games with small number of agents. However, in games with large populations and independently distributed types, the realized frequency of a type is identical to the ex-ante probability that an agent will draw this type[2].

Therefore, in all these cases, our mechanism requires *no additional information* than what is usually assumed. In fact, this argument holds even if there is correlation in the distribution of types (for example, if types are drawn from an ergodic Markov chain), as long as the drawing of types leads to a unique realized frequency. To make this point as clear as possible, in section 4.4 we re-examine the well-known economies by Spence (1973) and by Rothschild-Stiglitz (1976) and we show how our mechanism can be used in order to provide first-best allocations. To the best of our knowledge, this is the only mechanism that can solve these allocation problems efficiently[3].

[2]There are many papers which fall in this category, particularly in macroeconomics. For example, the literature on the implementation of the Walrasian correspondence in economies with adverse selection has traditionally made these assumptions. See also the papers by Prescott and Townsend (1984), Gale (1992 and 1996), Dubey and Geanakoplos (2002), Dubey, Geanakoplos and Shubik (2005), Bisin and Gottardi (2006), Rustichini and Siconolfi (2008).

[3]The Jackson-Sonnenschein mechanism requires infinite replicas of these economies in order to provide first-best allocations, while our mechanism can achieve the same outcome even if there exists

Moreover, the common knowledge of the realized frequency is assumed for our results because we consider general social choice rules. If we focus on the implementation of specific allocations on the Pareto frontier so that allocations depend only on ones type, we can implement the first-best as a unique equilibrium even if agents have heterogeneous beliefs or no information at all about the realized frequency (e.g. the Walrasian correspondence in the Rothschild-Stiglitz model). Our mechanism can still implement the desirable allocations truthfully, given that the social planner knows the realized frequency of types. This is because, as becomes clear in section 4, players' best-response correspondences depend on their beliefs about how many misreports will be detected by the mechanism and not on their ability to detect other agents' lies. For instance, this formulation fits the example of the store manager we provided earlier. The manager does not have to post the list of orders as we suggested earlier (though it was useful for the purposes of the exposition). It is sufficient that agents know that he knows them.

Furthermore, our results are fairly general. Types can be multi-dimensional, valuations can be independent or interdependent and the joint probability distribution over type-profiles allows for correlation across types or dependencies on the identity of the agents (different agents may face different probability distributions over types). The only restriction we impose on our notion of (Pareto) efficiency is Anonymity. Anonymity requires that the allocation, which an agent receives, depends only on his type (and possibly on the realized frequency of types) but not on his identity. It is a reasonable assumption which is satisfied by the majority of social choice rules. For instance, in many mechanism design papers, a mechanism is efficient if it implements the utilitarian social choice rule, which satisfies our definition of Anonymity[4].

We also provide necessary and sufficient conditions for full implementation. It should be stressed that we obtain our equilibrium by using iterated elimination of strictly dominated strategies and, hence, it is also a Bayes-Nash equilibrium. This contrasts with most of the existing papers, where the Bayesian equilibrium concept is used. Finally, we examine issues of robustness to small perturbations regarding the knowledge of the realized frequency of types and issues of participation constraints.

2 Related Literature

Our paper is most closely related to papers that use information aggregation to implement first-best allocations in economies with asymmetric information. Thus, in terms of spirit and research questions, Jackson and Sonnenschein (2007) is the paper closest to ours. They consider a specific set of agents, who 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

only one economy.

[4]See for example the papers by Mezzetti (2004), Jackson and Sonnenschein (2007).

the Bayes-Nash equilibria of these mechanisms converge to first-best allocations.

Our model differs from that of Jackson and Sonnenschein in four dimensions: i) we do not require multiple games to be played at the same time but we impose a stronger assumption on what is common knowledge (or, in certain cases, what is known by the central planner). ii) We allow for interdependent values, while they consider an independent values setting. iii) We allow for a more general joint probability over type profiles, since types can be independently or interdependently distributed in our formulation, and apart from preferences, types may concern other individual characteristics as well (productivity parameters, proneness to accidents, etc.). iv) We also allow for a more general social choice rule. In terms of results, if values are interdependent (but still independently distributed), the Jackson-Sonnenschein mechanism may have multiple equilibria in the limit, while we prove the uniqueness of the equilibrium under small perturbations.

McLean and Postlewaite (2002, 2004) also 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.

There are two main differences between their setting and ours. First, in the model of McLean and Postlewaite when private signals are perfectly correlated with the state of the world 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 can, at most, know the realized frequency of types (when the signal is perfect)[5]. Second, McLean and Postlewaite implement allocations arbitrarily close to first-best while we achieve exact first-best implementation even when agents face a slight uncertainty about the frequency of types, i.e. when private signals are slightly noisy.

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 large transfers which may violate ex-post feasibility. Also, Perry and Reny require the single crossing property on preferences which is a stronger restriction than ours. Our general framework can encompass auction design problems as well. Furthermore, our main focus is the uniqueness of the equilibrium, an issue which is not studied in these papers.

It is also noteworthy that in the framework of auction design the papers by Maskin (1992), Dasgupta and Maskin (2000) and Jehiel and Moldovanu (2001) show, in increasing generality, that efficiency and incentive compatibility can not be simultaneously

[5]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. As we have already mentioned, a small noise about the precision of these signals does not alter our results.

satisfied if the single crossing condition is violated or if signals are multidimensional. In that respect, the additional information of our environment allows us to overcome this impossibility and implement efficient outcomes, even if conditions, which are necessary in the standard mechanism design literature for implementation, are violated.

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 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.

3 The Economy

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 (so ϑ_i is the type of a single agent i). The vector $\boldsymbol{\theta}$ contains I elements and is a type-profile, a realization of a type for each agent. Each agent has private information about his own type, but does not know the types of the other agents. $\phi(\boldsymbol{\theta})$ is the ex-ante probability that the type-profile $\boldsymbol{\theta}$ will realize and $\Phi(\boldsymbol{\Theta})$ is the ex-ante probability distribution over all type-profiles.

S is the finite set of all states. Each state s is a complete description of the publicly available information. Depending on the application, this may include agents' features or public shocks. The probability distribution over states Π is a function of the type-profile $\boldsymbol{\theta}$. Therefore, $\pi(s|\boldsymbol{\theta})$ is the probability of state s arising, conditional on the type-profile $\boldsymbol{\theta}$.

β is the vector of realized relative frequencies of types in the population. Therefore, $\beta(\vartheta) = \lambda(\vartheta)/I$, where $\lambda(\vartheta)$ is the number of agents who have type ϑ in the population. It follows that, if one knows β, one can compute the exact number of agents, who have the same type in the population: $\lambda(\vartheta)$. The vector λ contains the exact number of agents for each type. We will often make allocations and mechanisms conditional on either λ or β, depending which one is most convenient in terms of exposition, but the equivalence between the two should be obvious. We will also use the terms relative frequencies, realized frequencies or simply frequencies interchangeably and we always refer to β. $\boldsymbol{\Theta}(\beta)$ is the set of all type-profiles consistent with the realized frequencies β, while $\Theta(\beta)$ is the collection of types which have realized, as can be inferred from β.

The above elements characterize the economy: $E = \{I, \Theta, \Phi, S, \Pi, \beta\}$. We assume that E is common knowledge. Given E, let $A(E)$ (or simply A) be the set of all feasible allocations, with elements $a \in A \subseteq R_+^{I\times S\times L}$, with $L \times S \geqslant 2$. Therefore, $a = \{a_1, .., a_i, ..., a_I\}$, where a_i is the individual allocation of agent i: $a_i \in R_+^{S\times L}$. L can be interpreted as the number of commodities in the economy. Each a is an S-tuple of feasible state-contingent allocations. In other words, the collection of feasible allocations may depend on the state of the world. Furthermore, we assume that preferences are represented by expected utility functions:

$$U_i(a_i) = \sum_{\theta_{-i}} \left[\sum_{s\in S} u_i(a_i, s)\ \pi\left(s|\vartheta_i, \boldsymbol{\theta}_{-i}\right) \right] \phi(\boldsymbol{\theta}_{-i}|\vartheta_i, \beta)\ , \quad \boldsymbol{\theta}_{-i} \in \boldsymbol{\Theta}_{-i}(\beta|\vartheta_i)$$

$U_i(a_i)$ is the expected utility to agent i when he receives allocation a_i, with $u_i(a_i, s)$ the decision-outcome payoff in state s (preferences may be state-dependent) and $\boldsymbol{\theta}_{-i}$ is a type-profile for all agents, excluding i, which is consistent with β[6]. Hence, $\phi(\boldsymbol{\theta}_{-i}|\vartheta_i, \beta)$ is the probability of the type profile $\boldsymbol{\theta}_{-i}$ realizing for the other agents, conditional on ϑ_i and β, and $\boldsymbol{\Theta}_{-i}(\beta|\vartheta_i)$ is the set of all type profiles for the other agents which is consistent with β, conditional on ϑ_i.

The formulation of the economy allows for modeling a wide variety of economic situations. Types may or may not be independently distributed, and the characteristics of agents may or may not depend on the types of other agents. Hence, both adverse-selection problems with independent or inter-dependent valuations can be seen as special cases of our formulation.

Additional Notation

In many cases, the analysis will require additional terms. In order to facilitate exposition, we would like to introduce as many of them as possible here, even though we leave the definition of some terms for later.

We have already defined $U_i(a_i)$. In many cases, however, it will be more convenient to examine preferences on different allocations by using the usual binary preference relation $\succ_i$: $a \succ_i b \Leftrightarrow U_i(a) > U_i(b)$, $a \succsim_i b \Leftrightarrow U_i(a) \geqslant U_i(b)$. Also, later on we will condition individual allocations on types, so that agents who have the same type receive the same final allocation. Then, it will be convenient to drop the subscript i and introduce the subscript ϑ instead. So, if agent i is of type ϑ and all agents of this type are to consume some allocation a_ϑ, then $U_i(a_\vartheta) = U_\vartheta(a_\vartheta)$. In other words, $U_\vartheta(a_\vartheta)$ is used to denote the expected utility of any agent of type ϑ. The notations $\succ_\vartheta, \succsim_\vartheta$ are interpreted in the same way.

The following definitions are also useful. $L_i(a_j)$ is the lower-contour set of agent i associated with individual allocation a_j: $L_i(a_j) = \{c \in R_+^{S\times L} : U_i(c) < U_i(a_j)\}$. $V_i(a_j)$

[6]Therefore, we implicit require the standard six axioms for expected utility representation: Completeness, Transitivity, Local Non-Satiation, Convexity, Continuity and Independence of Irrelevant Alternatives.

is the upper-contour set of agent i associated with a_j: $V_i(a_j) = \{c \in R_+^{S\times L} : U_i(c) > U_i(a_j)\}$. $L_\vartheta(a_{\vartheta'})$ and $V_\vartheta(a_{\vartheta'})$ are defined accordingly as above. $C_{i\epsilon}(a_j) = \{c \in R_+^{S\times L} : U_i(c) = U_i(a_j), \|c - a_j\| < \epsilon\}$ is the indifference plane of i in the neighborhood of a_j. $\underline{A}(a_i) = \{c \in R_+^{S\times L} : c_{ls} \leqslant a_{ils}, \forall\, ls\}$ is the set of individual allocations strictly less than a_i.

We have already defined $\lambda(\vartheta)$ as the number of agents with type ϑ according to β. In a mechanism, where agents send messages about their types, $\lambda_{\mathbf{m}}(\vartheta)$ is the number of agents who report type ϑ, where $\mathbf{m}$ denotes the message profile sent by agents: $\mathbf{m} = \{m_1, ..., m_i, ..., m_I\}$.

In some cases we denote explicitly that individual allocations are conditional on some other elements, like the message profile, other agents types or the realized frequencies of types. This helps to distinguish allocations while keeping the same symbol (a) to denote allocations throughout the paper. For example, the notation $a(\vartheta_i, \beta)$ is used in page 11 in order to show that, under Anonymity, Pareto efficient individual allocations depend only on one's type and the realized frequencies of types, but not one's identity. Similarly, $a_i(m_i, \mathbf{m}_{-i})$ in page 17 denotes that i's final allocation depends on both his report m_i and the reports of all other agents $\mathbf{m}_{-i}$. We clarify any change of notation on individual allocations every time we introduce a new one.

In other cases we perform simple mathematical operations between a scalar c and an individual allocation a_i, like addition ($a_i + c$) or multiplication (ca_i). In all cases involved, the interpretation is that we perform the same mathematical operation to all the state contingent commodities in a_i. Therefore, $a^*_\vartheta - \epsilon$ in page 17 implies that quantity ϵ is subtracted from all state contingent commodities, while $\frac{\lambda(\vartheta)}{\lambda_{\mathbf{m}}(\vartheta)} a^*_\vartheta$ in page 18 implies that all state contingent commodities in a^*_ϑ are scaled by $\frac{\lambda(\vartheta)}{\lambda_{\mathbf{m}}(\vartheta)}$. Finally, in some cases many agents receive identical individual allocations. The notation $n \times a_i$ denotes a multi-agent allocation which consists of n replicas of the individual allocation a_i and helps us to separate this case from the case of multiplication of an allocation by a scalar.

Summary of Notation

β:	realized relative frequencies of types
a_i:	individual allocation
a:	collective allocation
A:	set of feasible collective allocations
ϑ:	type
Θ:	set of types
$\boldsymbol{\theta}$:	type-profile
$\boldsymbol{\theta}_{-i}$:	type-profile of all agents except for i
$\mathbf{\Theta}$:	set of type-profiles
$\mathbf{\Theta}_{-i}(\beta\|\vartheta_i)$:	set of type-profiles excluding i, which are consistent with i's information set
$\phi(\boldsymbol{\theta})$:	ex-ante probability of type-profile $\boldsymbol{\theta}$

$\Phi(\boldsymbol{\Theta})$:	ex-ante probability distribution over type-profiles
$\phi(\boldsymbol{\theta}_{-i}\vert\vartheta_i,\beta)$:	ex-ante probability of $\boldsymbol{\theta}_{-i}$, conditional on i's information set
s:	state
S:	set of states
$\pi(s\vert\boldsymbol{\theta})$:	probability of state s conditional on type-profile $\boldsymbol{\theta}$
$u_i(a_i,s)$:	utility of agent i conditional on a_i and s
$U_i(a_i)$:	expected utility of allocation a_ϑ for agent i
$U_\vartheta(a_i)$:	expected utility of allocation a_ϑ for type ϑ
$\succ_i$:	preference relation of agent i
$\succ_\vartheta$:	preference relation of type ϑ
$L_i(a_j)$:	lower-contour set of allocation a_j for agent i
$V_i(a_j)$:	upper-contour set of allocation a_j for agent i
$C_{i\epsilon}(a_j)$:	indifference plane of i in the neighborhood of allocation a_j
$L_\vartheta(a_j)$:	lower-contour set of allocation a_j for type ϑ
$V_\vartheta(a_j)$:	upper-contour set of allocation a_j for type ϑ
$C_{\vartheta\epsilon}(a_j)$:	indifference plane of type ϑ in the neighborhood of allocation a_j
$\underline{A}(a_i)$:	set of individual allocations, which are less than a_i for at least one state-contingent commodity
m_i:	message of agent i
$\mathbf{m}$:	message-profile
$\mathbf{m}_{-i}$:	message-profile of all agents except for i
$\lambda(\vartheta)$:	number of agents with type ϑ
$\lambda_{\mathbf{m}}(\vartheta)$:	number of agents who have reported type ϑ according to message-profile $\mathbf{m}$
$a(\vartheta_i,\beta)$:	individual allocation conditional on ϑ_i and β
$a_i(m_i,\mathbf{m}_{-i})$:	individual allocation conditional on message-profile
$a_i \pm c$:	adding or subtracting c from all state-contingent commodities in a_i
ca_i:	multiplying all state-contingent commodities in a_i by c
$n \times a_i$:	number of times individual allocation a_i is provided

4 Implementation of First Best Allocations

4.1 Implementation

In this subsection we show that the conditions specified in section 3 are sufficient for the implementation of truthful strategies. Full implementation (i.e. the uniqueness of the truthful equilibrium) requires additional conditions, which we specify in subsections 4.2 and 4.3. The main idea is simple. The knowledge of the realized frequencies of types allows the construction of a direct mechanism, which provides allocations conditional on the message profile being consistent with these frequencies. If the message profile is not consistent with β, this is considered as an indication of lying by some agent, in which case the mechanism provides a "punishment" allocation. As a result, an agent reveals his information truthfully, if all other agents reveal their information truthfully as well.

Let $a^* = (a_1^*, a_2^*, ..., a_i^*, ..., a_I^*)$ be a Pareto efficient allocation of the economy. Let a^m be an individual allocation such that $a_{ls}^m = \min\{a_{ils}^*\}$ for every $i \in I$ and for each state-contingent commodity ls. By construction, $I \times a^m$ is feasible. Consider the direct mechanism $M_0(g, a)$, $g : M \to A$, in which agents state their type. Agents receive allocations according to the following message profiles:

- If $\lambda(\vartheta) = \lambda_{\mathbf{m}}(\vartheta)$, $\forall\ \vartheta \in \Theta(\beta)$, then $a_i = a_i^*, \ \forall\ i \in I$.
- If $\lambda(\vartheta) \neq \lambda_{\mathbf{m}}(\vartheta)$ for at least one $\vartheta \in \Theta(\beta)$, then $a_i = a^m, \ \forall\ i \in I$.

Claim 1: M_0 has a truthful equilibrium.

Proof: Suppose $I - 1$ agents report truthfully. By Local Non-Satiation, $U_i(a_i^*) \geqslant U_i(a^m)$. Therefore, it is a best-response for agent i to report truthfully as well.

This demonstrates that, if β is common knowledge, then this is a sufficient condition for truthful implementation in general economic environments. In fact, implementation of the truthful equilibrium is possible even when there is a single state contingent commodity. Hence, the implementation of first-best allocations is possible in the most well-known models of adverse selection (Akerlof (1970), Spence (1973), Rothschild-Stiglitz (1976)) if one makes the additional assumption that the realized frequencies of types are known.

Even though this is a strong assumption, in subsection 4.6, we show that as the number of agents increases, β converges to the ex-ante distribution of types. Hence, the standard assumptions of the literature are sufficient for implementation of first-best allocations when the number of agents is sufficiently large[7].

4.2 Full Implementation

In this section we provide sufficient conditions for full implementation. We make three assumptions additional to section 3. We then present a series of Lemmata, which are used in the proof of the main Proposition, and provide the main claim of the paper: if the realized frequencies of types are common knowledge, preferences satisfy the Local Non-Common Indifference Property (LNCIP) and the social choice rule satisfies Pareto efficiency and Anonymity, then a mechanism exists that fully implements it. The assumptions required for this result are the following.

Assumption 1: The Social Choice Rule satisfies Anonymity.

[7]Actually, for our results to obtain we do not require that the realized frequencies converge to the ex ante distribution. We only need that they converge to a unique distribution, given the correlation between draws.

Definition 1: A Social Choice Rule satisfies Anonymity if $a_i^* = a(\vartheta_i, \beta) = a_\vartheta^*,\ \forall i \in I$, where $\vartheta_i = \vartheta$.

Under Anonymity, agents who have identical types receive identical allocations. Therefore, an agent's identity per-se has no impact on the agent's final allocation. As a result, for any β there is a unique collection of allocations to be assigned to agents. The order of the allocations does depend on the type-profile $\boldsymbol{\theta}$, but the collection of individual allocations is the same for all type-profiles consistent with the same β.

It is also noteworthy that 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.

Assumption 2: Preferences satisfy the Local Non-Common Indifference Property (LNCIP).

This is a requirement that the intersection of the indifference planes around any individual allocation of any two agents with different types is of at least one dimension lower than the dimensions of the indifference planes themselves. In other words, if the indifference planes are n-dimensional (e.g. three-dimensional surfaces), the intersection around any allocation a_i is (n-1)-dimensional (e.g. curves). Formally:

Definition 2: Let $C_{i\epsilon}(a_j) = \{c \in R_+^{S\times L} : U_i(c) = U_i(a_j), \|c - a_j\| < \epsilon\}$. The **Local Non-Common Indifference Property** is satisfied if $\forall i \in I$, $\forall a_j \in R_+^{S\times L}$ and $\forall h \in I,\ \vartheta_h \neq \vartheta_i$, there exists $\bar{\epsilon}_{ih} > 0 : dim\,(C_{i\epsilon}(a_j) \cap C_{h\epsilon}(a_j)) \leqslant L \times S - 1\ ,\ \ \forall\, \epsilon < \bar{\epsilon}_{ih}$.

LNCIP is a weaker restriction than the Single-Crossing Property (SCP) which is usually used in the literature. For example, any pair of indifference curves that has finitely many intersections satisfies the LNCIP but it violates the SCP. Also, LNCIP allows for tangent indifference planes (as long as the tangent parts "miss" at least one dimension compared to the indifference planes), while the SCP does not. On the other hand, if SCP is satisfied then LNCIP is also satisfied[8]. Figure 1 provides two diagrams, which illustrate the LNCIP and distinguish it from the SCP.

[8]Note that we could alternatively characterize this restriction on preferences in terms of the axiomatic approach. Apart from the standard axioms (Completeness, Transitivity, Local Non-Satiation, Convexity, Continuity and Independence of Irrelevant Alternatives), we would require the Axiom of Local Non-Common Indifference. In this case, the only difference from the definition provided above is the definition of $C_{i\epsilon}(a_j)$: $C_{i\epsilon}(a_j) = \{c \in R_+^{S\times L} : c \sim_i a_j, \|c - a_j\| < \epsilon\}$.

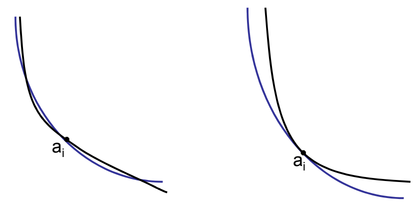

Figure 1: Indifference Curves satisfying LNCIP

Assumption 3: If for ϑ, ϑ' holds that $a^*_\vartheta \succ_{\vartheta'} a^*_{\vartheta'}$ and $a^*_{\vartheta'} \succ_\vartheta c$, $\forall c \in \underline{A}(a^*_\vartheta) \cap L_{\vartheta'}(a^*_{\vartheta'})$, then $\lambda(\vartheta') \geqslant \lambda(\vartheta)$, where ϑ and ϑ' are different types.

Assumption 3 ensures feasibility off-the-equilibrium-path and is used for deriving Lemma 4. Here, we provide the main economic intuition behind it with the help of figures 2 and 6. The main point is that given any two types ϑ, ϑ', such that ϑ' envies the first-best allocation of ϑ, one can find incentive compatible and feasible allocations for all agents of these two types.

To see why we need restrictions on the realized number of types, consider Figure 2 which depicts an economy with only two goods (X_1 and X_2) and two types. Suppose also that a^*_ϑ and $a^*_{\vartheta'}$ are the first-best allocations for types ϑ and ϑ' respectively. By construction, both types prefer a^*_ϑ to $a^*_{\vartheta'}$. In order to implement these allocations fully, we need to make sure that: (i) if ϑ' reports untruthfully, ϑ still prefers to report truthfully, (ii) if type ϑ reports truthfully, then type ϑ' prefers to report truthfully as well (these conditions are proved to be necessary for full implementation in subsection 4.3 and in the Appendix). The first condition requires to find some feasible allocation, say d, such that $d \succ_\vartheta a^*_{\vartheta'}$, while the second requires some c such that $a^*_{\vartheta'} \succ_{\vartheta'} c \succ_{\vartheta'} d$ (see also subsection 4.3). So the allocations $\{c, d\}$ can be seen as "punishment" for ϑ' for misreporting his type and as a "reward" for ϑ if he reports truthfully when the other type misreports.

However, off-the-equilibrium-path feasibility places restrictions on the pair $\{c, d\}$. By the definition of the Pareto frontier, the pair $\{a^*_\vartheta, a^*_{\vartheta'}\}$ is feasible, but, because we consider a very general economy, we do not know how the frontier behaves away from the pair of first-best allocations. This is not a problem when $\underline{A}(a^*_\vartheta) \cap L_\vartheta(a^*_{\vartheta'}) \subset \underline{A}(a^*_\vartheta) \cap L_{\vartheta'}(a^*_{\vartheta'})$, because then we can find $d \in \underline{A}(a^*_\vartheta)$ such that requirement (i) above is satisfied (this is shown in Figure 6 in the Appendix). In this case, whenever the number of types ϑ and ϑ' do not match the realized frequency of types, we let agents choose allocations from the set $\{a^*_{\vartheta'}, d\}$. Because these allocations are incentive compatible by construction, all ϑ choose d and all ϑ' choose $a^*_{\vartheta'}$. Furthermore, this is always feasible, because we can provide $\lambda(\vartheta)$ allocations d (as many as the number of agents who would like to receive them) and $\lambda(\vartheta')$ allocations $a^*_{\vartheta'}$ (again as many as the agents who would like to receive them). In other words, $\lambda(\vartheta) \times d$ and $\lambda(\vartheta') \times a^*_{\vartheta'}$ are both feasible due to the Pareto efficiency of $\{a^*_{\vartheta}, a^*_{\vartheta'}\}$.

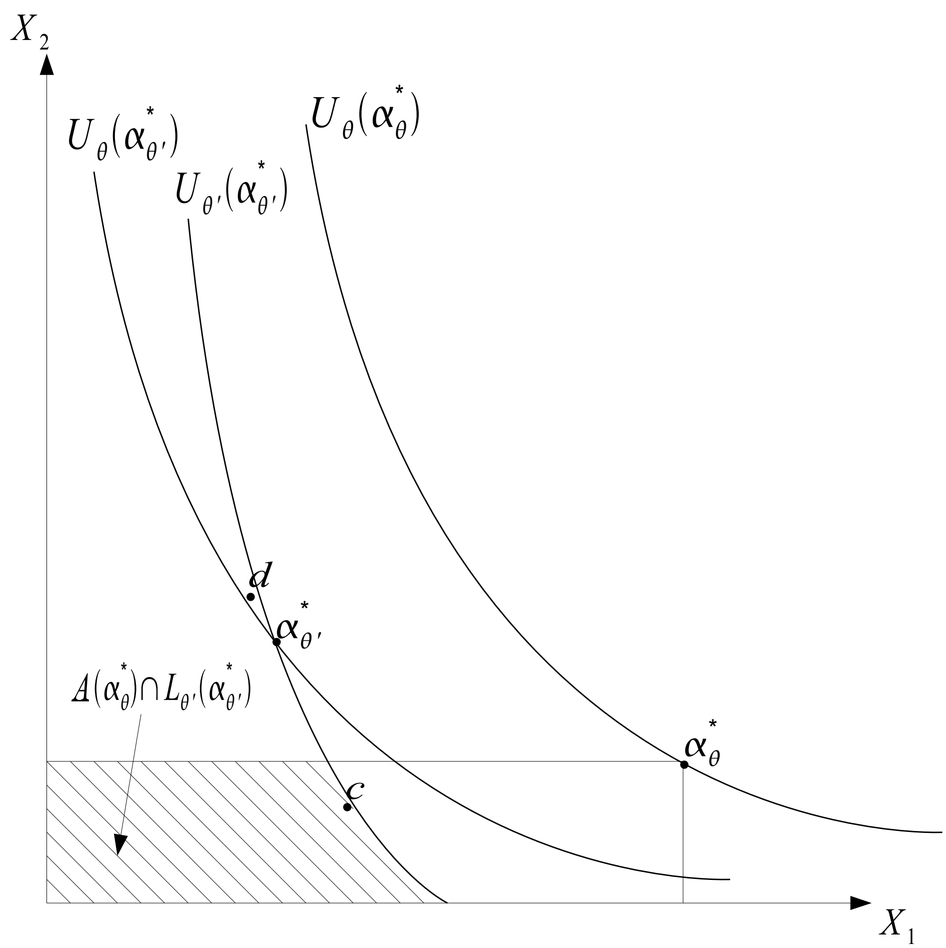

Figure 2: Feasible and Incentive Compatible Allocations. Case (ii):$\underline{A}(a^*_{\vartheta}) \cap L_{\vartheta'}(a^*_{\vartheta'}) \subset L_{\vartheta}(a^*_{\vartheta'})$

The problem arises when $\underline{A}(a^*_{\vartheta}) \cap L_{\vartheta'}(a^*_{\vartheta'}) \subset L_{\vartheta}(a^*_{\vartheta'})$, which is depicted in figure 2. In this case, if d is chosen in the interior of $\underline{A}(a^*_{\vartheta})$ it will either violate restriction (i) or (ii). As a result, d is chosen in the neighborhood of $a^*_{\vartheta'}$, while c is chosen in the interior of $L_{\vartheta}(a^*_{\vartheta'}) \cap \underline{A}(a^*_{\vartheta})$, as shown in the figure. The problem is that d, which is incentive compatible for type ϑ, is available only as many times as the number of ϑ' ($\lambda(\vartheta')$), while there are $\lambda(\vartheta)$ agents, who would like to receive this allocation off-the-equilibrium path. This means that, if $\lambda(\vartheta) > \lambda(\vartheta')$, some agents of type ϑ, will have to receive allocation c, which violates restriction (i). In this case, some ϑ types do not have an incentive to remain truthful if they believe that some ϑ' may lie and this leads to issues of multiplicity of equilibria. Assumption 3 is sufficient (but not necessary) for avoiding these problems[9].

Another sufficient condition is to assume the single-crossing property on the utility functions of different types, so that the types who are envied have steeper indifference planes than those types who envy them. Under this assumption, the case of figure 2 would not arise, so that Assumption 3 is not needed for the feasibility of the out-of-equilibrium-path allocations. In fact, the single-crossing property is a special case of our conditions, which are weaker sufficient conditions. In 4.3 we also provide the necessary conditions for full implementation.

We proceed by providing three results which hold for any Pareto efficient allocation. The combination of these results shows that every allocation on the Pareto frontier of an economy generates a "social ranking" among the agents of the economy, such that agents of "lower ranks" envy the allocations of "higher ranks". We exploit the common knowledge of this ranking, due to the common knowledge of β and the efficiency of the allocation, in order to construct a mechanism, which has a unique equilibrium and in which agents reveal their private information truthfully.

Lemma 1: Let PF(E) be the Pareto Frontier of economy E. Then, for every allocation a on the Pareto Frontier, there exists at least one agent $i \in I$, who does not envy the allocation of any other agent: $U_i(a_i) \geqslant U_i(a_j), \forall j \in I$.

Proof: See the Appendix

Lemma 2: For every allocation a on the Pareto Frontier, there exists at least one agent $i \in I$, whose allocation is not envied by any other agent: $U_j(a_j) \geqslant U_j(a_i), \forall j \in I$.

Proof: See the Appendix

Corollary 1: If $a \in PF(E)$, then Lemma 1 and 2 hold for any subset of I. Namely, let $\check{I} \subseteq I$ and let $\check{A} = \{a_i : i \in \check{I}\}$. Then, if $a \in PF(E)$, Lemma 1 and 2 hold for $\check{I}$ with regard to $\check{A}$ as well.

Proof: See the Appendix

[9]Note that we do not care about the opposite situation, where $\lambda(\vartheta') > \lambda(\vartheta)$, because then we have sufficient number of allocations for all the ϑ agents, and the fact that there are not enough allocations for the ϑ' types only increases their "punishment" for misreporting.

Lemma 1 and 2 provide two necessary conditions for Pareto efficiency. If these conditions are violated, then an allocation can not be Pareto efficient. However, they are not sufficient. One can easily find examples, where these conditions hold but the allocation is not on the Pareto frontier of the economy. Most importantly for our purposes, they imply that any Pareto efficient allocation exhibits a social ranking between groups of agents who envy and groups who are envied.

Let **Rank(K)** $= \{i \in I : U_i(a_i) \geqslant U_i(a_j), \forall j \in I\}$, be the set of agents who do not envy the allocation of any other agent. By Lemma 1, we know that this set is non-empty. Then, by removing this set of agents from the set I and applying Corollary 1, we can define **Rank(K-1)** $= \{i \in I - Rank(K) : U_i(a_i) \geqslant U_i(a_j), \forall j \in I - Rank(K)\}$. By iteration, we can define K groups, $1 \leqslant K \leqslant I$, such that the agents in each one of them do not envy any of the agents in their own group or groups with lower rank, but they envy the allocation of some agent(s) in groups with higher rank[10]. We will also refer to group **Rank(K)** as the group with the **highest rank** and group **Rank(1)** as the group with the **lowest rank**. Some additional results required for the proof of the main result come from the LNCIP and are provided in Lemma 3 and Lemma 4.

Lemma 3: Under LNCIP, there exists $\bar{\epsilon}(a_i)$ and a_ϑ, $a_{\vartheta'}$, such that:

(i) $a_\vartheta \succ_\vartheta a_i$, $\|a_\vartheta - a_i\| < \epsilon$
(ii) $a_{\vartheta'} \succ_{\vartheta'} a_i$, $\|a_{\vartheta'} - a_i\| < \epsilon$
(iii) $a_\vartheta \succ_\vartheta a_{\vartheta'}$, $a_{\vartheta'} \succ_{\vartheta'} a_\vartheta$

$\forall\ \epsilon < \bar{\epsilon}(a_i)$, $\forall\ \vartheta, \vartheta' \in \Theta$, $\vartheta \neq \vartheta'$, and $\forall\ a_i \in R_+^{S\times L}$.

Proof: See the Appendix

In effect, Lemma 3 states that, if the LNCIP holds, then in the neighborhood of any individual allocation a_i, there exists a set of allocations such that each agent of a certain type prefers a particular allocation over the rest. In other words, it is possible to find incentive compatible allocations for any type in the neighborhood of any allocation, which implies that it is possible to satisfy no-envy, at least in a local sense.

Lemma 4: Suppose $a^* \in PF(E)$ and Assumptions 1 and 2 hold. $\forall \vartheta, \vartheta' \in \Theta(\beta)$ there exist some feasible individual allocations $\{a_1(\vartheta, \vartheta'), a_2(\vartheta, \vartheta')\}$, such that, if $a^*_\vartheta \succ_{\vartheta'} a^*_{\vartheta'}$, then $a_1(\vartheta, \vartheta') \succ_\vartheta a^*_{\vartheta'} \succsim_\vartheta a_2(\vartheta, \vartheta')$, $a^*_{\vartheta'} \succsim_{\vartheta'} a_2(\vartheta, \vartheta') \succ_{\vartheta'} a_1(\vartheta, \vartheta')$.

[10]One extreme case is when an allocation exhibits no-envy, in which case Rank(K) contains the whole set of agents and Lemma 1 and 2 apply for all (egalitarian allocations). The other extreme case is when each rank-group contains a single agent, in which case the agents form a complete hierarchy, from the one who is envied by all the other agents to the one who is not envied by anyone else.

Proof: See the Appendix

First, in terms of notation, $a_1(\vartheta, \vartheta')$, $a_2(\vartheta, \vartheta')$ are used to denote that the allocations a_1 and a_2 depend on the types ϑ and ϑ'. This is useful because in Proposition 1 we use $\{a_1(\vartheta, \vartheta')$, $a_2(\vartheta, \vartheta')\}$ as off-the-equilibrium-path allocations if there is some discrepancy in the reported number of two types from the number expected under the realized frequencies. Obviously, then, the off-the-equilibrium-path allocations should depend on the types involved.

Second, in terms of meaning, Lemma 4 states that for any pair of types, such that one type envies the first-best allocation of the other, one can find a pair of feasible individual allocations such that the envied type prefers one of the two allocations over the other and over the first-best allocation of the other type. While the type, who envies, prefers his first-best allocation over these two allocations. This Lemma allows us to construct an off-the-equilibrium-path credible "reward" for ϑ if he reports truthfully even if ϑ' misreports, while at the same time these allocations turn into a "punishment" for ϑ' if he misreports when ϑ reports truthfully. In Proposition 1, we exploit Lemma 4 in conjunction with the ranking of agents on the Pareto frontier in order to prove the uniqueness of equilibrium of our mechanism.

This also brings us back to the discussion of Assumption 3. In Lemma 4, feasibility is ensured under the implicit assumption that the number of agents is equal across types. If this is not true, then additional restrictions on the number of realized types are required, specifically for the case where $\underline{A}(a^*_\vartheta) \cap L_{\vartheta'}(a^*_{\vartheta'}) \subset L_\vartheta(a^*_{\vartheta'})$. As we have already discussed, this is the case where the envied type ϑ, receives an allocation in the neighborhood of $a^*_{\vartheta'}$ off-the-equilibrium-path. Incentive compatibility is ensured whenever all agents of type ϑ (equal to $\lambda(\vartheta)$) receive this allocation with certainty, but feasibility implies that there are only $\lambda(\vartheta')$ such allocations available. Assumption 3 is sufficient for incentive compatibility and off-the-equilibrium-path feasibility to be satisfied simultaneously.

Lemmas 3 and 4, along with the knowledge of the "social ranking" of the allocations, allows us to construct a mechanism which makes it a dominant strategy for agents of higher rank to report their type truthfully. The main idea is that, if the number of agents, who report a specific type is higher than the number who have this type, according to the realized frequencies, then they all receive an allocation, which the "true" types prefer to the first-best allocations of the misreporting types, but the other types do not prefer. This acts as an effective punishment for lies by those who envy allocations of other types. Hence, we use iterated elimination of dominated strategies to prove the uniqueness of the proposed equilibrium.

A final note, before presenting our result. The mechanism, which we use in the proof, may induce a sub-game if the number of reports do not match the realized frequencies for two types. The game is one where the agents of the two misreported types choose an allocation from a "pool" of feasible and incentive compatible individual allocations. More specifically, if two types, ϑ and ϑ' are misreported, then the agents who reported these two types are sequentially drawn at random to choose an allocation from a collection of allocations. The collection contains $\lambda(\vartheta)$ times an identical allocation, which is incentive compatible for type ϑ, and $\lambda(\vartheta')$ times an allocation, which is incentive compatible for type ϑ'. Each time an agent chooses an allocation, this allocation is removed from the collection and the next agent chooses from the remaining allocations. Since this sub-game is induced when only two types are misreported, the number of agents, who are involved, is exactly equal to the number of individual allocations of the collection.

Formally, this game is represented by $G(\vartheta, \vartheta', \lambda(\vartheta) \times a_\vartheta, \lambda(\vartheta') \times a_{\vartheta'})$, where ϑ and ϑ' are the types involved and $\lambda(\vartheta) \times a_\vartheta$ denotes the number of times ($\lambda(\vartheta)$) of the individual allocation (a_ϑ), which is incentive compatible for ϑ: $a_\vartheta \succ_\vartheta a_{\vartheta'}$. It is easy to check that the unique sub-game perfect equilibrium of this game is for each agent to receive his most preferred allocation. With this in mind and with the use of Lemmas 1-4, we are now in position to present our main result.

Proposition 1: Assume that the economy E, described in section 3, satisfies Assumptions 2 and 3. Then, for every allocation $a^* \in PF(E)$, which satisfies Assumption 1, there exists a mechanism, for which a^* is the unique Bayes-Nash equilibrium allocation and agents report their private information truthfully.

Proof: The proof is done by construction. Let $a^* \in PF(E)$, which satisfies Anonymity, and let $a^*(\boldsymbol{\theta})$ be the first-best allocation which is to be implemented for each type-profile. Also, let $\hat{a}_\vartheta(a, \epsilon)$ denote an individual allocation in the ϵ-neighborhood of allocation a which is incentive compatible for type ϑ, in the sense of Lemma 3[11], and let $a_1(\vartheta, \vartheta'), a_2(\vartheta, \vartheta')$ be individual allocations as constructed by Lemma 4. Recall that $\lambda(\vartheta)$ and $\lambda_{\mathbf{m}}(\vartheta)$ is the number of agents of type ϑ according to β and the received messages $\mathbf{m}$, respectively, and a^m is the minimum allocation, as defined in 4.1.

Each agent reports his type m_i and a final allocation is received according to the following mechanism $M_1(g, a)$:

i) If $\mathbf{m} \in \boldsymbol{\Theta}(\beta)$, then $a_i(m_i, \mathbf{m}_{-i}) = a^*_{m_i}$, $\forall i \in I$.

ii) If $\mathbf{m}$ is such that for only two types, (ϑ, ϑ'), the number of reported agents is different from the number of agents according to the realized frequencies by one, specifically $\lambda_{\mathbf{m}}(\vartheta) = \lambda(\vartheta) + 1$, $\lambda_{\mathbf{m}}(\vartheta') = \lambda(\vartheta') - 1$, then:

- If $a^*_\vartheta \succ_\vartheta a^*_{\vartheta'}$, $a^*_{\vartheta'} \succ_{\vartheta'} a^*_\vartheta$, then, for the agents who reported types ϑ, ϑ', the mechanism induces $G(\vartheta, \vartheta', \lambda(\vartheta) \times (a^*_\vartheta - \epsilon), \lambda(\vartheta') \times (a^*_{\vartheta'} - \epsilon))$[12].

[11]So, ϑ prefers $\hat{a}_\vartheta(a, \epsilon)$ to a and to any other allocation $\hat{a}_{\vartheta'}(a, \epsilon)$, which is provided for any other type in the neighborhood of a.

[12]ϵ is strictly positive for all state-contingent commodities and it is sufficiently small so that $a^*_\vartheta - \epsilon \succ_\vartheta$

- If $a^*_\vartheta \succ_{\vartheta'} a^*_{\vartheta'}$, then, for the agents who reported types ϑ, ϑ', the mechanism induces $G(\vartheta, \vartheta', \lambda(\vartheta) \times a_1(\vartheta, \vartheta'), \lambda(\vartheta') \times a_2(\vartheta, \vartheta'))$.
- If $a^*_{\vartheta'} \succ_\vartheta a^*_\vartheta$, agents who report type ϑ' receive allocation $a^*_{\vartheta'}$ and agents who report type ϑ receive allocation $\frac{\lambda(\vartheta)}{\lambda_{\mathbf{m}}(\vartheta)} a^*_\vartheta$.
- For all $m_k \neq \{\vartheta, \vartheta'\}$, $a_k(m_k, \mathbf{m}_{-i}) = a^*_{m_k}$.

iii) For any other case, $a_i(\vartheta, \mathbf{m}_{-i}) = \hat{a}_\vartheta(a^m, \epsilon)$, $\forall \vartheta \in \Theta$.

Under the mechanism above, it is a strictly dominant strategy for all agents with types of rank(K) to report their type truthfully. To see this consider the different beliefs of an agent of rank(K) (say i of type ϑ) about the messages that other agents will send. If i believes that all other agents will report their type truthfully, then the best-response for him is to report truthfully. This is because $a^*_\vartheta \succ_i a^*_\vartheta - \epsilon$, in the case he reports another type, who does not envy a^*_ϑ, and $a^*_\vartheta \succ_i \frac{\lambda(\vartheta')}{\lambda_{\mathbf{m}}(\vartheta')} a^*_{\vartheta'}$, in the case he reports a type, who envies a^*_ϑ (recall that a rank(K) agent does not envy anyone.).

If i believes that only one other agent will misreport, then i still prefers to report his type truthfully, irrespectively of who misreports. Say that i believes that j is of different type (say ϑ'), does not envy a^*_ϑ and that j will misrepresent her preferences as being of type ϑ. If i reports that he is of type ϑ', then the two lies will cover each other and i receives $a^*_{\vartheta'}$. But if he reports ϑ, then $\lambda_{\mathbf{m}}(\vartheta') = \lambda(\vartheta') - 1$ and $\lambda_{\mathbf{m}}(\vartheta) = \lambda(\vartheta) + 1$. In the latter case, $G(\vartheta, \vartheta', \lambda(\vartheta) \times (a^*_\vartheta - \epsilon), \lambda(\vartheta') \times (a^*_{\vartheta'} - \epsilon))$ is induced and i receives $a^*_\vartheta - \epsilon$. Since $a^*_\vartheta - \epsilon$ is constructed to be strictly preferred by i to $a^*_{\vartheta'}$, i strictly prefers to report truthfully.

The same argument holds if i believes that j is of type ϑ', that j envies a^*_ϑ and that j will report ϑ. Note that, by construction of the set $\{a_1(\vartheta, \vartheta'), a_2(\vartheta, \vartheta')\}$ (see also Lemma 4), there are $\lambda(\vartheta)$ times which allocation $a_1(\vartheta, \vartheta')$ is feasible and $\lambda(\vartheta')$ times which allocation $a_2(\vartheta, \vartheta')$ is feasible. Since i strictly prefers $a_1(\vartheta, \vartheta')$ to $a^*_{\vartheta'}$, he prefers to report truthfully. Also, note that whenever i believes that j misreports, i strictly prefers to report truthfully than to send any other message $\vartheta \neq \{\vartheta, \vartheta'\}$, because, in the latter case, i receives $\hat{a}_\vartheta(a^m, \epsilon)$, which makes him strictly worse-off.

In the case where i believes that multiple misrepresentations will take place, then, irrespectively of his message, $\mathbf{m} \neq \mathbf{\Theta}(\beta)$ (if all reports but one cancel out then we go back to the analysis of the previous cases). This means that his message, alone, can not hide the fact that some agent(s) misrepresents(misrepresent) her(their) type(s). His best response remains to report truthfully: $U_i(\hat{a}_\vartheta(a^m, \epsilon)) > U_i(\hat{a}_{\vartheta'}(a^m, \epsilon))$, $\forall \vartheta' \neq \vartheta$ (recall that $I \times a^m$ is feasible). We conclude that, under all possible beliefs, i strictly prefers to report truthfully.

Given this, it is a best response for an agent of rank(K-1) to report his type truthfully as well. Say that agent i, who is of rank(K-1) and type ϑ, envies the allocation of some type ϑ' of rank(K). Of course, if i believes that some agent of type ϑ' will report as being of type ϑ, then the best response for i is $m_i = \vartheta'$, but, as we showed, this cannot

$a^*_{\vartheta'}$ and $a^*_{\vartheta'} - \epsilon \succ_{\vartheta'} a^*_\vartheta$.

be an equilibrium[13]. Hence, if i believes that all agents will report truthfully, he prefers to report truthfully as well. If he believes that only one agent of the same or lower rank will misreport their types as his own, he will still prefer to reveal his type truthfully, for the same type of reasoning as in the case of an agent of rank(K). Finally, if he believes that many agents will misreport their types, he still prefers to receive an incentive compatible allocation (by construction) than misrepresenting his own type. Therefore, given that rank(K) agents report truthfully, agents of rank(K-1) also report truthfully.

By induction, we conclude that for an agent of Rank(κ), if all agents of higher rank are expected to report truthfully their types, his best-response is to report truthfully, irrespectively of the actions of agents of the same or lower rank. Since it is a dominant strategy for rank(K) agents to report truthfully, then, by iterated elimination of strictly dominated strategies, the only possible equilibrium is when all agents report truthfully. Therefore, the unique Bayes-Nash equilibrium of the mechanism is for all agents to reveal their type and to receive the allocation $a^*_{\vartheta_i}, \forall i \in I$. ■

The result depends crucially on the fact that the rank of types is known. This is due to the realized frequencies of types being common knowledge. On the other hand, Anonymity ensures that agents do not gain any strategic benefit from their personal identity. For instance, even if β is common knowledge, if different type-profiles result in different ranks between types, then it may not be a dominant strategy for any agent to reveal his type truthfully. As one's rank, in this case, also depends on the realized types of the other agents, there may be situations where an agent misreports his type in order to force someone to misreport as well. This may cause multiplicity of equilibria. In other words, if Anonymity fails, implementation is still possible, but full implementation may fail.

The LNCIP is also required for the uniqueness of the equilibrium, as it allows for agents to strictly improve their payoff if they report truthfully. Once again, if LNCIP is violated, then one can still construct mechanisms which implement the first-best allocations, but the uniqueness of the equilibrium may not be possible. Therefore, the common knowledge of the realized frequencies, Anonymity and LNCIP (along with Assumption 3) are jointly sufficient conditions for full implementation of first-best allocations, but they are not necessary.

We would also like to comment on the advantages of our mechanism in comparison to the existing literature (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. Fourth, full implementation is also achieved if the

[13] This argument also makes clear that our paper is not one of dominant strategy implementation, as only rank(K) individuals have dominant strategies.

equilibrium concept is changed to iterated elimination of strictly dominated strategies, which is, in fact, the solution concept we use in the proof of Proposition 1. Therefore, our mechanism is not limited only to Bayesian implementation.

Finally, Assumptions 1,2 and 3 are relatively weak and there are many cases of interest that comply with them. To demonstrate this, in 4.4, we provide some well-known examples of economies with hidden types and the solutions that our framework provides. But first, we characterize the problem by providing necessary and sufficient conditions for full implementation.

4.3 Full Implementation: Necessary and Sufficient Conditions

Condition 1: Suppose $a^* \in PF(E)$. $\forall \vartheta, \vartheta' \in \Theta(\beta)$ such that $a^*_{\vartheta} \succ_{\vartheta'} a^*_{\vartheta'}$, $\exists \ a \in A$ such that: (i) $a_{\vartheta} \succ_{\vartheta} a^*_{\vartheta'}$, and (ii) $a^*_{\vartheta'} \succ_{\vartheta'} a_{\vartheta}$.

Proposition 2: Condition 1 is necessary for full implementation.

Proof: Suppose that Condition 1 is not satisfied. Full implementation of a^* requires that $g(\mathbf{m}) = a^*$ if $m_i = \vartheta_i$, $\forall i \in I$ and that the strategy profile $m_i = \vartheta_i$, $\forall i \in I$ is the unique Bayes-Nash equilibrium. Consider any direct mechanism $M(g, a)$, which specifies some allocation $a(\mathbf{m}) \neq a^*$, whenever $\mathbf{m}$ is such that $\lambda_{\mathbf{m}}(\vartheta'') \neq \lambda(\vartheta'')$ for some $\vartheta'' \in \Theta(\beta)$ (whenever this is the case, then, by common knowledge of β, it follows that $m_i \neq \vartheta_i$ for some $i \in I$). Suppose that, apart from i (of type ϑ) and j (of type ϑ'), incentive compatibility is satisfied for all other agents and that they report truthfully (this is done in order to check the necessity of the condition).

Because Condition 1 is violated, then either part (i) or part (ii) of the condition is violated (or both). This means that at least one of the following will hold: (i) $a_i(m_i = \vartheta, m_j = \vartheta, \mathbf{m}_{-i,j}) \succ_j a^*_{\vartheta'}$, (ii) $a^*_{\vartheta'} \succ_i a_i(m_i = \vartheta, m_j = \vartheta, \mathbf{m}_{-i,j})$. In case (i), truthful reporting is not equilibrium, because, if everyone else reports truthfully, j's best-response is $m_j = \vartheta$ (incentive compatibility is violated for j). In case (ii), there are multiple equilibria because, if the truthful equilibrium exists, then so does another equilibrium, where i reports type ϑ' and j reports type ϑ. To see this, notice that if i believes that j is of type ϑ' and that $m_j = \vartheta$, then his best-response is $m_i = \vartheta'$, in which case it is also a best-response for j to report $m_j = \vartheta$. Finally, in the case where both parts of Condition 1 are violated, then there can be no truthful equilibrium (as j strictly prefers to report ϑ, if everyone else reports truthfully). In all cases, full implementation is impossible. ■

Condition 1 is similar in spirit to Bayesian Monotonicity, which is necessary for full implementation in economies with incomplete information (Jackson, 1991). In our case, full implementation is possible, if there is a feasible allocation through which some agent (i) "signals" cases of misreport. As a result, not all efficient allocations are fully implementable when β is common knowledge. Note that Condition 1 holds whenever the number of agents of lower-rank are less or equal to the number of agents of higher ranks. Assumption 3 in section 4.2 made this restriction clear. On the other hand, Condition 1 is weaker than Assumption 3, and may hold in cases where this assumption is violated.

Furthermore, Condition 1 is sufficient for full implementation if one allows for mechanisms with games that do not have an equilibrium in pure strategies (for example integer games, as in Maskin (1999) or modulo games, as in Jackson (1991))[14]. This is because one can rule out undesirable equilibria with multiple misrepresentations of types (sub-case (iii) in the mechanism of Proposition 1) by making agents to play such a game, whenever the message-profile is not consistent with β by more than one message. However, if one restricts attention to mechanisms where agents send only messages about their types, the following condition is also required.

Condition 2: Suppose $a^* \in PF(E)$. There exists allocation an $a \in A$, such that $a^*_\vartheta \succ_\vartheta a_\vartheta$ and $a_\vartheta \succ_\vartheta a_{\vartheta'}$ $\forall\, \vartheta, \vartheta' \in \Theta(\beta)$.

Condition 2 ensures that whenever there are more than one misrepresentations of types, it is a best-response for one of the "liars" to deviate and report truthfully, while it is not a best-response to deviate from truth-telling. It becomes apparent that Assumption 3 and the LNCIP satisfy Condition 1 (Lemma 4), while LNCIP also satisfies Condition 2 (Lemma 3). Jointly, Condition 1 and 2 are necessary and sufficient for full implementation for this restricted set of mechanisms[15].

4.4 Examples

Spence (1973)

The Spence economy consists of two types. Group I has low productivity $\underline{a}$ and its proportion in the population is q_1. Group II has high productivity $\overline{a}$ and its proportion in the population is $1-q_1$. Acquiring y units of education costs $y/\underline{a}$ for Group I and $y/\overline{a}$ for Group II. 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 low productivity. The education acquired by Group II is a deadweight loss, but necessary for credible signaling.

[14] See the Appendix for the proof. We omit it here, since it is similar to the proof of Proposition 1.

[15] See the Appendix for the proof.

Assume that the total population N is common knowledge. Then Nq_1 is the total number of agents of Group I and $N(1-q_1)$ is the total number of agents of Group II. Based on 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 Group I and II is Nq_1 and $N(1-q_1)$, respectively, then agents who report Group I receive wage $w_{G_I} = \underline{a}$ and those who report Group II, receive wage $w_{G_{II}} = \overline{a}$. Otherwise, those who report Group I receive $w_{G_I} = \underline{a}$ and those who report Group II, are asked to undertake one unit of education and receive $w_{G_{II}} = \underline{a} + \epsilon$, 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 an $\overline{a}$-type. It is clear that, irrespectively of the reports of the other agents, it is a dominant strategy for her to report $\overline{a}$, since $\overline{a} > \underline{a}$ and $\underline{a} + \epsilon - \frac{1}{\overline{a}} > \underline{a}$. Then, it is a best-response for an $\underline{a}$-type to report truthfully as well. This is because $\underline{a} > \underline{a} + \epsilon - \frac{1}{\underline{a}}$. Hence, all agents report truthfully in equilibrium and acquire zero education. In Figure 2 contract a_0 denotes the offer to high-productivity workers when lies are detected.

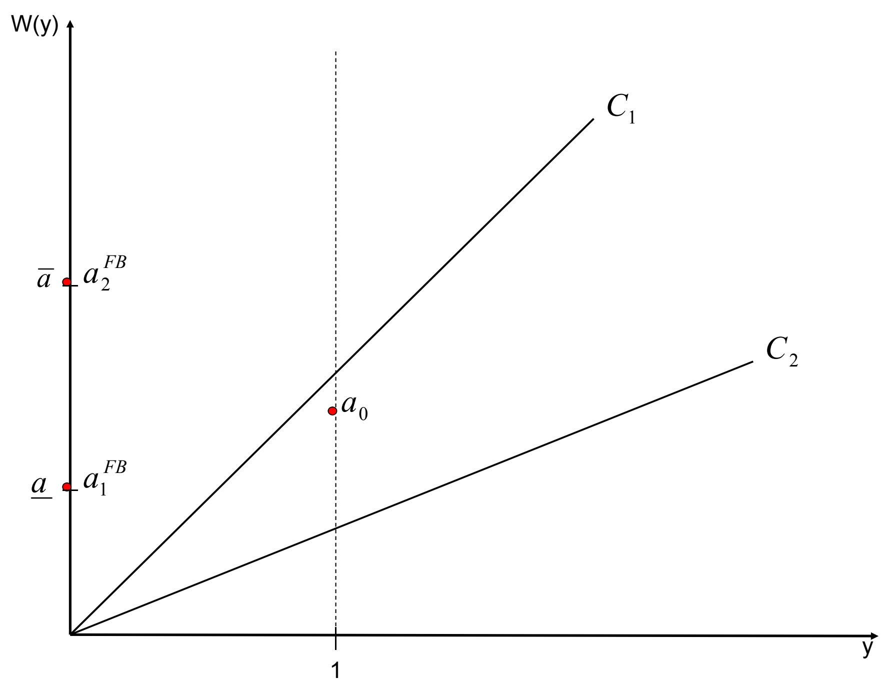

Figure 3: Spence, 1973

Rothschild-Stiglitz (1976)

Consider the Rothschild-Stiglitz economy. There is a continuum of risk-averse agents with mass equal to one and a risk-neutral entrepreneur. There is one commodity. Agents have a stochastic endowment with two possible states, H, L, with corresponding endowments w_H and w_L, where $w_H > w_L$. An agent's expected utility function depends on her consumption on both individual states: $U(c_L, c_H)$. The Bernoulli utility function $u(.)$ is strictly increasing and concave. There are two types of agents. A proportion λ_1 of the population are of type 1 and face a high probability of suffering from the low endowment state: p_H. The remaining $1 - \lambda_1$ are of type 2 and have a low probability of w_L: $p_L < p_H$. Types are private information, but the rest characteristics of the economy are common knowledge.

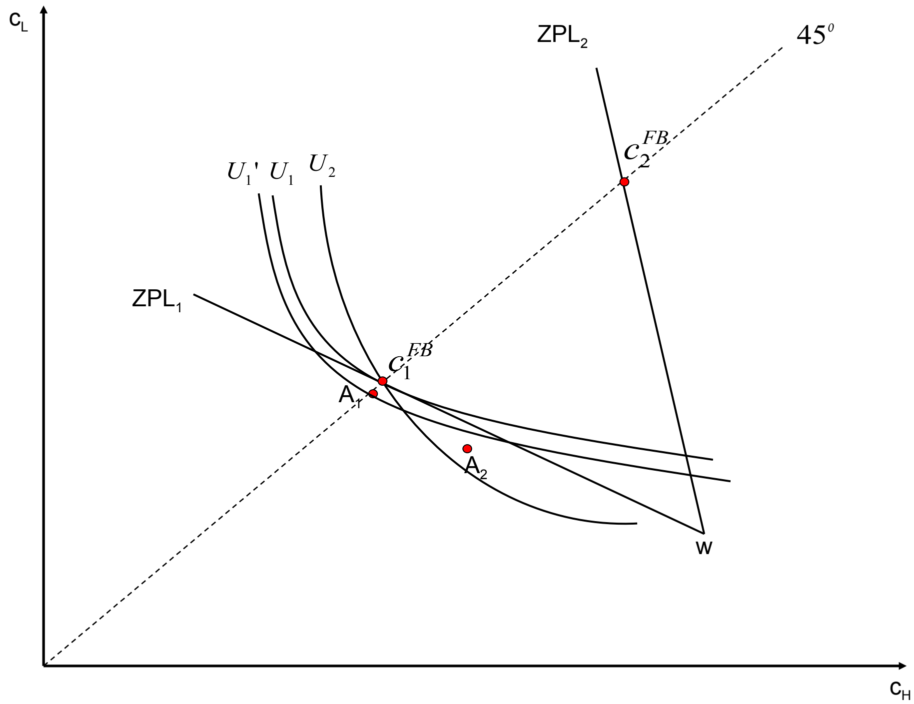

Figure 4: Rothschild-Stiglitz, 1976

Assuming that insurees have full bargaining power and hence the entrepreneur makes no profits from her services, the following mechanism can be utilized in order to implement first-best allocations (see also Figure 4). All agents report their type. If the message-profile matches the realized frequencies of types then each agent receives the insurance contract that corresponds to her message (C_1^{FB} and C_2^{FB} are the state-contingent allocations resulting from the first-best insurance contracts for 1 and 2 respectively). Otherwise, agents who report type 1, receive an insurance contract which results to allocation A_1, while agents who report type 2, receive A_2.

Notice that, by construction, $A_2 \succ_2 C_1^{FB} \succ_2 A_1$ and $C_1^{FB} \succ_1 A_1 \succ_1 A_2$. Also, providing any combination of these individual allocations to the agents of the economy is feasible, since they all lie in the interior of $\underline{A}(C_2^{FB})$. Therefore, Condition 1, is satisfied. It is easy to check that it is a dominant strategy for type 2 to report truthfully. Given this, it is a best-response for any agent of type 1 to report truthfully, as well. Therefore, the proposed mechanism has a unique Bayes-Nash equilibrium, which is truthful.

For this result to obtain, we have implicitly assumed that potential deviations from equilibrium take place by an ϵ mass of same-type agents so that the aggregate number of messages is affected by these deviations. Alternatively, one can use Al-Najjar (2004) to transform the economy into one with discrete large number of agents, which is equivalent to having a continuum of agents but preserves the results of discrete games. Therefore, our claim is that one can solve the allocation problem in the Rothschild-Stiglitz economy *without any additional information* than what was originally assumed. To the best of our knowledge, this is the only mechanism that achieves this result.

4.5 Robustness to Small Perturbations

So far we have assumed that the realized frequencies of types are commonly known with perfect precision. This is a very strong assumption, and hence we would like to make sure that small relaxations of it would not change our results dramatically. As it turns out, if there is a sufficiently small noise about β, then our main claim still holds.

Let Γ be the set of all possible frequencies of types that can be generated by $\boldsymbol{\Theta}$. By definition, $\bigcup_{\gamma\in\Gamma} \boldsymbol{\Theta}(\gamma) = \boldsymbol{\Theta}$. Suppose, now, that there is a small noise about the probability of the frequency β. Agents have a probability distribution over the set of γ. With probability $1 - \sum\limits_{\gamma\in\Gamma} \epsilon_\gamma$, the frequency of types β will be realized, while ϵ_γ is the probability that some other frequency γ will be realized, with $\epsilon_\gamma > 0, \forall\gamma \in \Gamma$.

We maintain the assumption that each agent knows his own type with certainty but has no information about the other agents' type. The expected utility of agent i has to be modified in order to include the uncertainty over frequencies:

$$U_i(a_i) = (1 - \sum_{\gamma\in\Gamma} \epsilon_\gamma) \sum_{\boldsymbol{\theta}_{-i}\in\boldsymbol{\Theta}_{-i}(\beta|\vartheta_i)} \left[\sum_{s\in S} u_i(a_i, s)\ \pi\left(s|\vartheta_i, \boldsymbol{\theta}_{-i}\right)\right] \phi(\boldsymbol{\theta}_{-i}|\vartheta_i, \beta)$$
$$+ \sum_{\gamma\in\Gamma} \epsilon_\gamma \left[\sum_{\boldsymbol{\theta}_{-i}\in\boldsymbol{\Theta}_{-i}(\gamma|\vartheta_i)} \left[\sum_{s\in S} u_i(a_i, s)\ \pi\left(s|\vartheta_i, \boldsymbol{\theta}_{-i}\right)\right] \phi(\boldsymbol{\theta}_{-i}|\vartheta_i, \gamma)\right]$$

We also assume that for every γ there is a Pareto optimal allocation to be implemented, which satisfies Anonymity: $a^*(\gamma)$, with individual allocations $a_i^*(\gamma) = a^*(\vartheta_i, \gamma)$.

In the case of uncertainty about the realized frequency, the rank of each agent is also uncertain, as different γ may correspond to different sets of realized types and different ranks. The problem then would be similar to the problem when the Anonymity property is violated. However, if this uncertainty is sufficiently small, the equilibrium strategies of agents will not change. To see this, consider an agent i who has the highest rank under β (and potentially other ranks for other γ's). If he knows that β is the realized frequency with certainty, then under the mechanism presented in 4.2, he would strictly prefer to report his type truthfully than report any other type:

$$U_i(\vartheta_i, \mathbf{m}_{-i}|\beta) > U_i(\vartheta', \mathbf{m}_{-i}|\beta), \forall \vartheta' \neq \vartheta_i \in \Theta, \forall \mathbf{m}_{-i} \in M$$

Adding a small uncertainty about β means that his expected utility by reporting his type truthfully becomes:

$$U_i(\theta_i, \mathbf{m}_{-i}) = (1 - \sum_{\gamma \in \Gamma} \epsilon_\gamma) U_i(\vartheta_i, \mathbf{m}_{-i}|\beta) + \sum_{\gamma \in \Gamma} \epsilon_\gamma U_i(\vartheta_i, \mathbf{m}_{-i}|\gamma)$$

It is evident that, if ϵ_γ is sufficiently small for every γ, the expected utility of i approaches the expected utility under β and hence it remains a strictly dominant strategy to report his type truthfully. The argument can be repeated for any other agent j of different rank according to β. Given a sufficiently small vector of probabilities $\boldsymbol{\epsilon}$, j expects all higher-rank agents to report truthfully and his best-response is to report truthfully as well, irrespectively of the messages send by agents of the same or lower ranks. Hence, there exists some vector $\boldsymbol{\epsilon}$, with strictly positive elements, such that the equilibrium strategies under certainty over β remain the unique equilibrium strategies under uncertainty over β.

Corollary 2: If the realized frequencies of types are uncertain but there is a sufficiently high probability that some β will be realized, then the mechanism of Proposition 1 fully implements the first-best allocations for every realized frequency.

Proof: It follows from the analysis above.

It is noteworthy that, due to the fact that truthful revelation of one's type is the only equilibrium action for all agents, the desirable individual allocations will be implemented for any γ. In other words, the almost certainty about β makes agents to report their type truthfully irrespectively of the frequencies of types that are eventually realized. As a consequence, agents receive first-best allocations for all realized frequencies. This confirms that our result is robust to small perturbations of the information structure and it is not just a construction of perfect knowledge of β.

4.6 Convergence to Ex-Ante Distributions

So far we have shown our main result and that it is robust to small uncertainty about the realized frequencies. We also want to point out that if the number of agents becomes very large then the realized frequencies converge to the ex-ante probability distribution of types[16], in which case our informational assumptions converge to the widely used assumptions in the standard mechanism design literature, i.e. agents know the ex-ante probability of each type occurring. This allows us to relate our formulation and results to large economies with adverse selection problems, and make the claim that in these economies, because the realized frequencies are effectively common knowledge, one can implement first-best allocations.

Of course, this requires some restrictions on the joint probability function Φ. The easiest way is to assume that types are independently and identically distributed. This means that the probability of acquiring type ϑ, $\tau(\vartheta)$, is the same across all agents and the draws of types from the ex-ante distribution are uncorrelated. Then, by directly applying the Weak Law of Large Numbers we get:

$$\lim_{I\to\infty}\left(\frac{\lambda_{(\vartheta)}}{I}\right)=\tau(\vartheta)$$

This is exactly the information provided by β: the number of agents, for whom type ϑ has realized. Hence, at the limit, the relative frequency of types in the population coincides with the ex-ante probability distribution [17]. Hence, our mechanism can be applied to economies with large populations without requiring any additional information than the standard mechanism design literature on asymmetric information and with minimal restrictions on the joint probability function.

4.7 Participation Constraints

A final note is required regarding the issue of participation constraints. In many important applications of adverse selection problems, agents are given the opportunity not to participate in a contract or in a mechanism if the expected utility they anticipate by entering is less than some exogenously given threshold. In our model, however, we have completely ignored any participation constraint restrictions. Fortunately, this omission does not result in loss of generality. If participation constraints are to be taken into consideration, then this only restricts the points of the Pareto frontier that satisfy these constraints and does not alter the rest of the analysis[18].

[16]Note that, for our results to obtain we do not require that the realized frequencies converge to the ex ante distribution. We only need that they converge to a unique distribution, given the correlation between draws.

[17]Notice, however, that other formulations of the Law of Large Numbers do not require independently or identically distributed types. For example, suppose that the type generating process is an ergodic Markov chain. Then, as the number of draws becomes infinitely large, the empirical distribution of types converges to a unique distribution (see for example Grinstead and Snell, 1997). Clearly, in this case, draws may be correlated, but, as long as the mechanism designer knows the transition matrix of the Markov chain and assuming that all draws take place before the mechanism is played, then β can be estimated with arbitrary precision as the number of agents approaches infinity. Generally, our mechanism can be applied in all cases where the realized frequencies of types converge to a unique distribution as the population becomes very large.

[18]Of course, in all interesting problems, the intersection of all participation constraints with the Pareto-frontier is non-empty. Notice that, in off-the-equilibrium-path situations, the resulting allo-

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 Bayes-Nash equilibrium, where all agents reveal their type truthfully and they receive a first-best allocation. Our result relies on information aggregation and appropriately chosen punishments. If the realized frequencies of types are known (perfectly or imperfectly), 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 in terms of the aggregate information 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 allocations when lies are detected. We show that such punishments exist when the indifference curves of different types are not locally identical, meaning that in the neighborhood of any allocation one can find other allocations such that each type prefers one of these over the rest.

It should be stressed that we obtain our equilibrium by using iterated elimination of strictly dominated strategies and, hence, it is also a Bayes-Nash equilibrium. This contrasts with most of the existing papers, where the Bayesian equilibrium concept is used. Furthermore, the assumption on the realized frequencies of types being common knowledge is needed because we consider general social choice rules. If we focus on the implementation of specific allocations on the Pareto frontier so that allocations depend only on one's type (and not on the realized frequencies of types, as we have examined in this paper), we can implement the first-best as a unique equilibrium even if agents have heterogeneous beliefs or no information at all about the realized frequencies of types. Our mechanism can still implement the desirable allocation truthfully, given that the social planner knows these frequencies. This is because players' best-response correspondences depend on their beliefs about how many misreports will be detected by the mechanism and not on their ability to detect other agents' lies. Finally, an interesting question is whether the implementation of first-best allocations in this setting can be achieved through a decentralized mechanism. We plan to address this question in the near future.

cations may violate certain participation constraints. But as long as agents decide and commit on their participation before the mechanism is played (based on the expectation of an outcome, which results from some equilibrium of the sub-game), then the uniqueness and efficiency of the equilibrium guarantees the participation of all agents.

References

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

Al-Najjar, N. (2004): "Aggregation and the Law of Large Numbers in Large Economies", Games and Economic Behavior, 47:1, 1-35.

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.

Bisin, A., and P. Gottardi (2006): "Efficient Competitive Equilibria with Adverse Selection", Journal of Political Economy, 114, 485-516.

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

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

Dasgupta, P., and E. Maskin (2000): "Efficient Auctions", Quarterly Journal of Economics, 115, 341-388.

Dubey, P., and J. Geanakoplos (2002): "Competitive Pooling: Rothschild-Stiglitz Reconsidered", Quarterly Journal of Economics, 117, 1529-1570.

Dubey, P., J. Geanakoplos and M. Shubik (2005): "Default and Punishment in General Equilibrium", Econometrica, 73, 1-37.

Gale, D. (1992): "A Walrasian Theory of Markets with Adverse Selection", Review of Economic Studies, 59, 229-255.

Gale, D. (1996): "Equilibra and Pareto Optima of Markets with Adverse Selection", Economic Theory, 7, 207-235.

Grinstead, C., and L. Snell (1997): "Introduction to Probability", American Mathematical Society

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.

Jehiel, P., and B. Moldovanu (2001): "Efficient Design with Interdependent Valuations", Econometrica, 69, 1237-1259.

Maskin, E. (1992): "Auctions and Privatization", in Privatization, ed. by H. Siebert. Institut für Weltwertschaft an der Universität Kiel.

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.

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.

Prescott, E., and R. Townsend (1984): "Pareto Optima and Competitive Equilibria with Adverse Selection and Moral Hazard", Econometrica, 52, 21-45.

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

Rustichini, A., and P. Siconolfi (2008): "General Equilibrium in Economies with Adverse Selection", Economic Theory, 37, 1-29.

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

Appendix

Lemma 1: Let PF(E) be the Pareto Frontier of economy E. Then, for every allocation a on the Pareto Frontier, there exists at least one agent $i \in I$, who does not envy the allocation of any other agent: $U_i(a_i) \geqslant U_i(a_j), \forall j \in I$.

Proof: Suppose that the claim does not hold. Then, all agents envy at least one other agent: $\forall\ a_i\ \exists\ j \in I, j \neq i : U_i(a_j) > U_i(a_i)$. But, since this holds for all agents, then there exists at least one reassignment of individual allocations among the I agents such that some of them are made strictly better-off and the rest remain as well-off as under a.

In order to find one such reassignment, use the following algorithm. Pick an arbitrary $i \in I$ and let $\bar{i} = \{j \in I : U_i(a_j) > U_i(a_i)\}$, be the set of agents whom i envies. Reassign a_j, for some $j \in \bar{i}$, to i. If $i \in \bar{j}$, then reassign a_i to j and stop the reassignment. If $i \notin \bar{j}$, then reassign some a_h, $h \in \bar{j}$ to j and then proceed to agent h. Continue until you reach some agent k, such that either $i \in \bar{k}$ or there exists some $l \in \bar{k}$, whose allocation a_l has already being reassigned. In the first case, reassign allocation a_i to k and stop the reassignments. In the latter case, ignore all reassignments preceding agent l (these agents retain their original allocations), reassign to l the allocation a_k and stop the reassignments.

Since the set of agents is finite and all agents envy at least one allocation, after at most I reassignments, the algorithm above will end-up in some agent, whose allocation has already been reassigned, or the first agent, where the reassignment started. In this case, a reassignment of allocations has been found, which makes some agents in I better-off (from agent l until agent k) while the rest remain equally well-off. This constitutes a Pareto improvement and violates the initial assumption that $a \in PF(E)$. ■

Lemma 2: For every allocation a on the Pareto Frontier, there exists at least one agent $i \in I$, whose allocation is not envied by any other agent: $U_j(a_j) \geqslant U_j(a_i), \forall j \in I$.

Proof: The proof is similar to the proof of Lemma 1. Suppose that the claim does not hold. Then, all agents are envied by at least one other agent: $\forall a_i\ \exists\ j \in I, j \neq i : U_j(a_i) > U_j(a_j)$. But, this implies that there exists at least one reassignment of individual allocations among the I agents such that some of them are made strictly better-off and the rest remain as well-off as under a.

In order to find one such reassignment, use the following algorithm. Pick an arbitrary $i \in I$ and reassign a_i to one of the agents in the set $\underline{i} = \{j \in I : U_j(a_i) > U_j(a_j)\}$. Then reassign a_j. If $i \in \underline{j}$, then reassign a_j to i and stop the reassignment. If $i \notin \underline{j}$, then reassign a_j to some arbitrary $h \in \underline{j}$ and repeat the reassignment. Continue until you reach some agent k, such that there exists some $l \in \underline{k}$, whose allocation a_l has already being reassigned. Ignore all reassignments preceding agent l (these agents retain their original allocations), reassign to l the allocation a_k and stop the reassignments.

Since the set of agents is finite and all allocations are envied by at least one agent, after at most I reassignments, the algorithm above will end-up in some agent whose allocation has already been reassigned. In this case, we have found a reassignment of allocations which makes some agents in I better-off while the rest remain equally well-off. This constitutes a Pareto improvement and violates the initial assumption that $a \in PF(E)$. ■

Corollary 1: If $a \in PF(E)$, then Lemma 1 and 2 hold for any subset of I. Namely, let $\check{I} \subseteq I$ and let $\check{A} = \{a_i : i \in \check{I}\}$. Then, if $a \in PF(E)$, Lemma 1 and 2 hold for $\check{I}$ with regards to $\check{A}$ as well.

Proof: Take any subset of agents $\check{I}$ of the set I. Suppose that Lemma 1 and 2 do not hold over the set $\check{A}$, which is the set of individual allocations of the agents in $\check{I}$. Then, it is possible to find a reassignment of allocations between the agents in $\check{I}$, such that some of them will be made better-off while the rest remain as well-off. But that is a Pareto-improvement for some agents in I, which contradicts the assumption that $a \in PF(E)$. ■

Lemma 3: Under LNCIP, there exists $\overline{\epsilon}(a_i)$ and a_ϑ, $a_{\vartheta'}$, such that:

(i) $a_\vartheta \succ_\vartheta a_i$, $\|a_\vartheta - a_i\| < \epsilon$
(ii) $a_{\vartheta'} \succ_{\vartheta'} a_i$, $\|a_{\vartheta'} - a_i\| < \epsilon$
(iii) $a_\vartheta \succ_\vartheta a_{\vartheta'}$, $a_{\vartheta'} \succ_{\vartheta'} a_\vartheta$

$\forall\ \epsilon < \overline{\epsilon}(a_i)$, $\forall\ \vartheta, \vartheta' \in \Theta$, $\vartheta \neq \vartheta'$, and $\forall\ a_i \in R_+^{S\times L}$.

Proof: Recall that $C_{i\epsilon}(a_j) = \{c \in R_+^{S\times L} : U_i(c) = U_i(a_j), \|c - a_j\| < \epsilon\}$. Also, recall that $L_j(a_i)$ is the lower-contour set of agent j associated with allocation a_i and $V_j(a_i)$ is the upper-contour set.

H is a $L \times S - 1$ hyper-plane, which passes through a_i, and is perpendicular to the marginal rate of substitution of some type's indifference plane, which also passes through a_i. **H** splits the space of allocations in two sub-spaces, A_1 and A_2. In each of these sub-spaces, and due to the LNCIP, there exists some $\overline{\epsilon} > 0$ such that for every $\epsilon < \overline{\epsilon}$, within the open ball $B_\epsilon(a_i)$, the upper contour set of a type is a subset of the upper contour set of some other type (see also the figure below).

Say that agent k is the type with the smallest upper contour set within ball $B_\epsilon(a_i)$

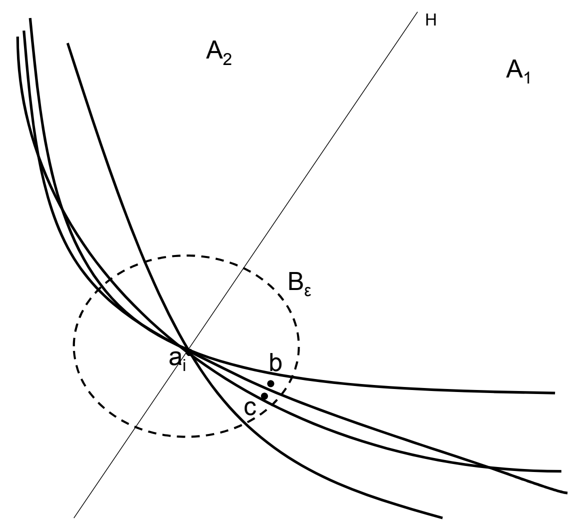

Figure 5: LNCIP and Local Incentive Compatibility

and subspace A_1: $V_k(a_i)\bigcap B_\epsilon(a_i)\bigcap A_1 \subset V_l(a_i)\bigcap B_\epsilon(a_i)\bigcap A_1, \forall l \in \Theta$. Then, by LNCIP, there exists some allocation $b \in B_\epsilon(a_i)$ such that a_i is strictly preferred to b by agents of type k, but the agents of all other types strictly prefer b to a_i: $b \in L_k(a_i)$ and $b \in V_l(a_i), \forall l \in \Theta$.

Likewise, there exists allocation c, which does not belong in the two smallest upper contour sets within $B_\epsilon(a_i)$ but it is within all the other upper contour sets, which means that a_i is strictly preferred by type k to b and c, b is strictly preferred by the type with the second smallest contour set to a_i and c and all the other types prefer c to a_i and b. By induction, one can construct $\Theta - 1$ allocations in the ϵ-neighborhood of a_i, such that the agents of one type strictly prefer one allocation over all the other. The properties described by Lemma 3 follow immediately. ■

Lemma 4: Suppose $a^* \in PF(E)$ and Assumptions 1 and 2 hold. $\forall \vartheta, \vartheta' \in \Theta(\beta)$ there exist some feasible individual allocations $\{a_1(\vartheta, \vartheta'), a_2(\vartheta, \vartheta')\}$, such that, if $a^*_\vartheta \succ_{\vartheta'} a^*_{\vartheta'}$, then $a_1(\vartheta, \vartheta') \succ_\vartheta a^*_{\vartheta'} \succsim_\vartheta a_2(\vartheta, \vartheta')$, $a^*_{\vartheta'} \succsim_{\vartheta'} a_2(\vartheta, \vartheta') \succ_{\vartheta'} a_1(\vartheta, \vartheta')$.

Proof: Because a Pareto efficient allocation is feasible by definition, any allocation $c \in \underline{A}(a^*_\vartheta) \cup \underline{A}(a^*_{\vartheta'})$ is feasible. Also, due to the Pareto efficiency of a^* and the fact that ϑ' envies the first-best allocation of ϑ, $L_{\vartheta'}(a^*_{\vartheta'}) \cap \underline{A}(a^*_\vartheta) \neq \emptyset$. Take an individual allocation c inside this intersection and arbitrarily close to the indifference plane of ϑ'

that passes through $a^*_{\vartheta'}$. Therefore, $a^*_{\vartheta'} \succ_{\vartheta'} c$. There are two possible sub-cases to consider, which are shown in figures 6 and 2 respectively.

Case (i): $c \succ_{\vartheta} a^*_{\vartheta'}$. In this case, let $a_1(\vartheta, \vartheta') = c$ and $a_2(\vartheta, \vartheta') = a^*_{\vartheta'}$ and this completes the proof. $\lambda(\vartheta)$ allocations c and $\lambda(\vartheta')$ allocations $a^*_{\vartheta'}$ are feasible on aggregate.

Case (ii): $a^*_{\vartheta'} \succ_{\vartheta} c$. In this case, by LNCIP, it is possible to find an allocation d very close to $a^*_{\vartheta'}$ such that: $d \succ_{\vartheta} a^*_{\vartheta'}$ and $c \succ_{\vartheta'} d$. Because c is in the interior of $\underline{A}(a^*_{\vartheta})$, it is always possible to find such points (we could define distance ϵ and make sure that $B_\epsilon(c) \cap U_{\vartheta'}(a^*_{\vartheta'}) \neq \emptyset$, while $B_\epsilon(d) \cap U_{\vartheta'}(a^*_{\vartheta'}) = \emptyset$, where $B_\epsilon(c)$ is the open ball with radius ϵ around c). Therefore, let $a_1(\vartheta, \vartheta') = d$ and $a_2(\vartheta, \vartheta') = c$. $\lambda(\vartheta)$ allocations d and $\lambda(\vartheta')$ allocations c are feasible on aggregate. ■

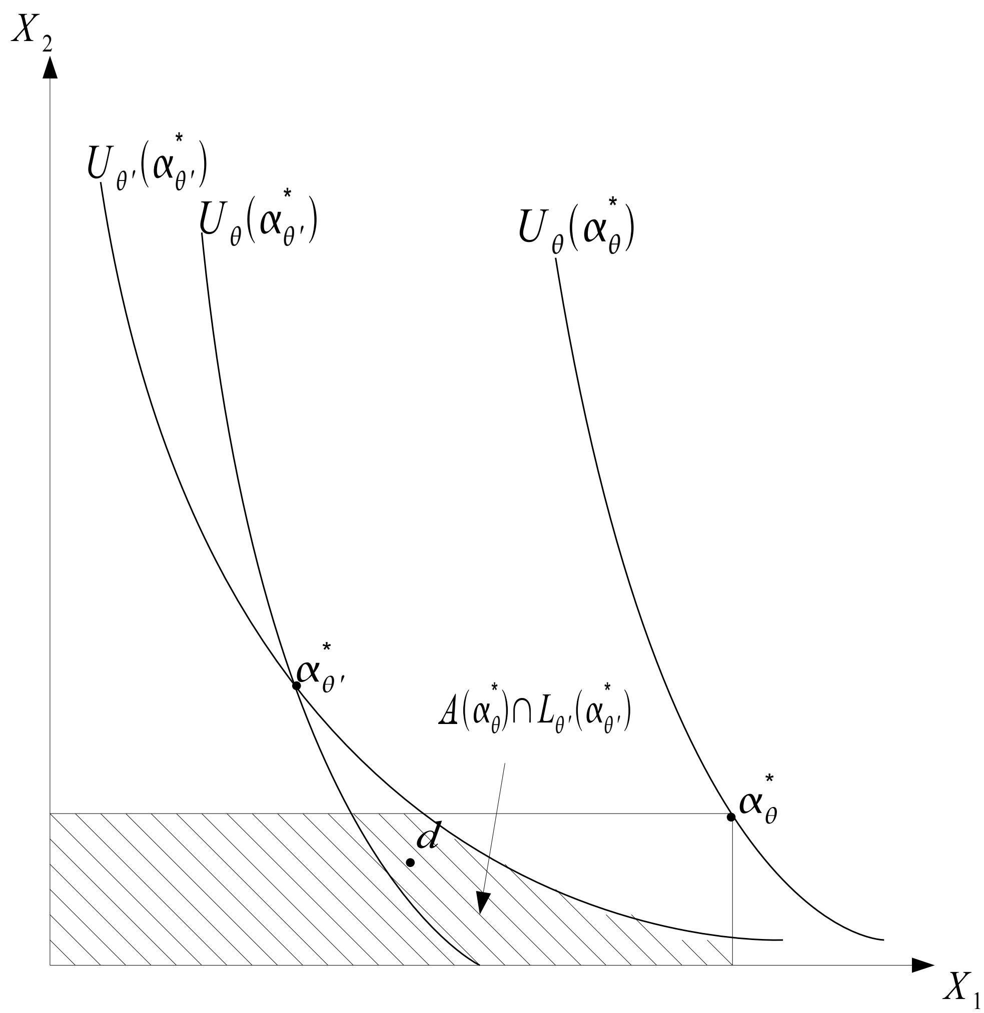

Figure 6: Feasible and Incentive Compatible Allocations. Case (i): $\underline{A}(a^*_{\vartheta}) \cap L_{\vartheta}(a^*_{\vartheta'}) \subset \underline{A}(a^*_{\vartheta}) \cap L_{\vartheta'}(a^*_{\vartheta'})$

Proposition 3: In the space of mechanisms, which permit sub-games with no equilibrium in pure strategies, Condition 1 is sufficient for full implementation.

Proof: Suppose that $G_I(P, A)$ is a simultaneous move game $G : P \to A$ with I players, and assume that G has no Nash equilibrium in pure strategies (examples include Jackson (1991) and Maskin (1999)). Also, arbitrarily restrict the payoffs of G_I such that the maximum possible payoff for any type is lower than if he were to receive the first-best allocation of any other type[19]. Let $R_i(p_{-i}, G)$ be the best-response correspondence of agent i if game G_I is played. Finally, suppose that Condition 1 is satisfied and thatβ is common knowledge. The mechanism below fully implements any Pareto efficient allocation which satisfies Anonymity.

Each agent reports his type m_i and a final allocation is received according to the following mechanism $M(g, a)$:

i) If $\mathbf{m} \in \mathbf{\Theta}(\beta)$, then $a_i(m_i, \mathbf{m}_{-i}) = a^*_{m_i}$, $\forall i \in I$.

ii) If $\mathbf{m}$ is such that for only two types, $(\vartheta,\ \vartheta')$, the number of reported agents is different from the number of agents according to the realized frequencies by one, specifically $\lambda_{\mathbf{m}}(\vartheta) = \lambda(\vartheta) + 1$, $\lambda_{\mathbf{m}}(\vartheta') = \lambda(\vartheta') - 1$, then:

- If $a^*_\vartheta \succ_\vartheta a^*_{\vartheta'}$, $a^*_{\vartheta'} \succ_{\vartheta'} a^*_\vartheta$, then, for the agents who reported types ϑ, ϑ', the mechanism induces $G(\vartheta, \vartheta', \lambda(\vartheta) \times (a^*_\vartheta - \epsilon), \lambda(\vartheta') \times (a^*_{\vartheta'} - \epsilon))$.
- If $a^*_\vartheta \succ_{\vartheta'} a^*_{\vartheta'}$, then, for the agents who reported types ϑ, ϑ', the mechanism induces $G(\vartheta, \vartheta', \lambda(\vartheta) \times a_1(\vartheta, \vartheta'), \lambda(\vartheta') \times a_2(\vartheta, \vartheta'))$.
- If $a^*_{\vartheta'} \succ_\vartheta a^*_\vartheta$, agents who report type ϑ' receive allocation $a^*_{\vartheta'}$ and agents who report type ϑ receive allocation $\frac{\lambda(\vartheta)}{\lambda_{\mathbf{m}}(\vartheta)} a^*_\vartheta$.
- For all $m_k \neq \{\vartheta, \vartheta'\}$, $a_k(m_k, \mathbf{m}_{-i}) = a^*_{m_k}$.

iii) For any other case, the mechanism induces G_I.

If more than one misreport is detected, M induces G_I, which has no equilibrium[20]. Therefore, there can be no equilibrium of the mechanism where agents believe that more than two misreports will be detected. Conditional on that, it is a strictly dominant strategy for the agents of the highest rank to report truthfully their type. To see this, take agent i of type ϑ and suppose that his rank is K. Agent i's only possible equilibrium beliefs are that: either (i) all other agents will report truthfully or (ii) one other agent will misreport or (iii) there will be multiple misreports but they will cover

[19] An easy way to do this is to multiple all payoffs of G_I with an arbitrarily small but positive number.

[20] More than one misreport detected means that either $\lambda_\beta(\vartheta) \neq \lambda_{\mathbf{m}}(\vartheta)$ for more than two types or that $\lambda_{\mathbf{m}}(\vartheta) - \lambda_\beta(\vartheta) \geqslant 2$ and $\lambda_{\mathbf{m}}(\vartheta') - \lambda_\beta(\vartheta') \leqslant 2$ for some types ϑ, ϑ'.

each other (e.g. type ϑ_k reporting as type ϑ_l and vice versa) apart from one. Case (ii) and (iii) are strategically equivalent for i as his response induces the same allocation.

If i believes that all other agents will report their type truthfully then his best response is to report truthfully as well. Otherwise, he receives either the allocation $a^*_\vartheta - \epsilon$ or the allocation $\frac{\lambda(\vartheta')}{\lambda_{\mathbf{m}}(\vartheta)} a^*_{\vartheta'}$. Clearly, i strictly prefers a^*_ϑ to the above allocations and his best response is to report his type truthfully.

If, on the other hand, i believes that an agent (say j of type ϑ') of rank(K) will misreport his type as ϑ, then i, by reporting truthfully, receives $a^*_\vartheta - \epsilon$, while by reporting type ϑ' he receives $a^*_{\vartheta'}$. By construction, $a^*_\vartheta - \epsilon \succ_\vartheta a^*_{\vartheta'}$. If i reports any other type then his payoff will be even less due to the restrictions on the payoffs of G_I. Hence, i's best response is to report truthfully.

If i believes that an agent (say m of type ϑ'') of a lower rank (and who envies a^*_ϑ) will misreport his type as ϑ, then a similar argument goes through. Reporting truthfully is strictly preferred to reporting any other type, since $a_1(\vartheta, \vartheta'') \succ_\vartheta a^*_{\vartheta''}$. Finally, if i believes that some agent m of type ϑ_m will misreport his type to ϑ_n, then i prefers reporting truthfully and receiving a^*_ϑ to reporting untruthfully and receiving some payoff induced by G_I.

Hence, for all beliefs that can be consistent with equilibrium, all agents of rank K strictly prefer to report their type truthfully. Given this and by following the same reasoning, agents of $rank(K-1)$ strictly prefer to report truthfully as well. By induction and iterated elimination of strictly dominated strategies, we conclude that all ranks will report truthfully and hence the unique Bayes-Nash equilibrium of the mechanism is for all agents to report their type truthfully. ■

Proposition 4: Condition 1 and 2 are jointly sufficient for full implementation.

Proof: Suppose that Condition 1 and 2 are satisfied and that β is common knowledge. The mechanism below fully implements any Pareto efficient allocation which satisfies Anonymity. Each agent reports his type m_i and a final allocation is received according to the following mechanism $M(g, a)$:

i) If $\mathbf{m} \in \mathbf{\Theta}(\beta)$, then $a_i(m_i, \mathbf{m}_{-i}) = a^*_{m_i}$, $\forall i \in I$.

ii) If $\mathbf{m}$ is such that for only two types, $(\vartheta,\ \vartheta')$, the number of reported agents is different from the number of agents according to the realized frequencies by one, specifically $\lambda_{\mathbf{m}}(\vartheta) = \lambda(\vartheta) + 1$, $\lambda_{\mathbf{m}}(\vartheta') = \lambda(\vartheta') - 1$, then:

- If $a^*_\vartheta \succ_\vartheta a^*_{\vartheta'}$, $a^*_{\vartheta'} \succ_{\vartheta'} a^*_\vartheta$, then, for the agents who reported types ϑ, ϑ', the mechanism induces $G(\vartheta, \vartheta', \lambda(\vartheta') \times (a^*_\vartheta - \epsilon), \lambda(\vartheta) \times (a^*_{\vartheta'} - \epsilon))$.
- If $a^*_\vartheta \succ_{\vartheta'} a^*_{\vartheta'}$, then, for the agents who reported types ϑ, ϑ', the mechanism induces $G(\vartheta, \vartheta', \lambda(\vartheta) \times a_1(\vartheta, \vartheta'), \lambda(\vartheta') \times a_2(\vartheta, \vartheta'))$.

- If $a^*_{\vartheta'} \succ_\vartheta a^*_\vartheta$, agents who report type ϑ' receive allocation $a^*_{\vartheta'}$ and agents who report type ϑ receive allocation $\frac{\lambda(\vartheta)}{\lambda_{\mathbf{m}}(\vartheta)} a^*_\vartheta$.
- For all $m_k \neq \{\vartheta, \vartheta'\}$, $a_k(m_k, \mathbf{m}_{-i}) = a^*_{m_k}$.

iii) For any other case, an allocation $\tilde{a}$, which satisfies Condition 2, is implemented.

The mechanism above is identical to the mechanism of Proposition 3, with the only exception that, if more than one misreport is detected, then instead of inducing a game without an equilibrium, the mechanism provides an allocation which is constructed according to Condition 2. By construction of $\tilde{a}$, all types prefer to report truthfully if they believe that many misreports will be detected.

Therefore, even if a rank(K)-agent believes that there will be several detections of misreports, he still prefers to report truthfully. He also prefers to report truthfully than reporting any other type, if he believes that there is only one misreport (say reporting some type ϑ' or ϑ'' when some ϑ' misreports), because $a_1(\vartheta, \vartheta') \succ_\vartheta a^*_{\vartheta'} \succ_\vartheta \tilde{a}_{\vartheta''}$. Since his best-response remains the same for all other beliefs, this means that any agent of rank(K) has a strictly dominant strategy to report truthfully. Therefore, by following the same reasoning as in the proof of Proposition 3 and by iterated elimination of strictly dominated strategies, we conclude that the mechanism has a unique Bayes-Nash equilibrium, at which all agents report their type truthfully. ■