

Information Aggregation and Adverse Selection*

Aristotelis Boukouras[†]

Kostas Koufopoulos[‡]

This Version: 30 May 2010

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 interim distribution of types is common knowledge (the exact number of agents for each type), then a mechanism exists, which is consistent with truthful revelation of private information and which implements first-best allocations of resources as the unique Bayes-Nash equilibrium. Our result requires weak restrictions on preferences (Local Incentive Compatibility Property) and on the Pareto correspondence (Anonymity) and it is robust for small enough noise around the interim distribution. Our paper is useful in understanding the power of information aggregation in alleviating incentive constraints and is particularly pertinent to games with large populations, in which case the interim distribution of types approaches the ex-ante distribution.

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, Thodoris Diasakos, Christopher Hennessy, Matthew Jackson, Eric Maskin, Claudio Mezzetti, Asako Ohinata, Motty Perry, Philip Reny, Paul Youdell and Nikos Zygouras, for their very useful comments and remarks, for which we feel grateful.

[†]University of Warwick, Department of Economics, a.boukouras@warwick.ac.uk

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

1 Introduction

As first shown by the papers of Akerlof (1970), Spence (1973) and Rothschild and Stiglitz (1976), hidden-types (adverse selection) problems can have significant consequences in terms of efficiency on economic outcomes¹. 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 out 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 (what we call the interim distribution of types), then it is possible, under fairly general conditions, to implement first-best allocations. More precisely, we consider an economy with asymmetric information and finite agents, each one of whom has private information about his type. We also assume that i) the interim-distribution of types is common knowledge, ii) preferences satisfy the Local Incentive Compatibility Property and iii) the social choice set satisfies Anonymity². Given these general conditions, we show that it is possible to construct a mechanism which has a unique Bayes-Nash equilibrium, where all agents truthfully reveal their type and they receive a first-best allocation.

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 informational shock. For example, a retail store has received pre-paid orders from its customers, has already the goods in stock and is ready to make the deliveries. However, the records on the orders get destroyed due to an accident and the store's manager does not know who made each order. What is he to do? Can he induce the customers to truthfully reveal the orders they have made without them making unreasonable claims or receiving orders that were meant for other customers? We claim that this is possible, as long as the

¹The 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 truthfully reveal their information 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.

²Since we are considering an economy of incomplete information, different realizations of types, which are consistent with the same interim-distribution, result in different desirable allocations. Therefore, we use the term Social Choice Set instead of the term Social Choice Rule or Correspondence, which usually refers to complete information environments. See also Jackson (1991) and Palfrey and Srivastava (1989).

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 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 million 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 interim distribution is known with perfect precision, we subsequently prove that it holds for 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 interim-distribution is known, then one can aggregate the messages that all agents are sending out and uncover any misreporting(s), even if the identity of the liar is not known. As a consequence, appropriately designed punishments for lying can induce agents to truthfully reveal their information.

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 should also point out that we derive this result for a general hidden-types environment. Types can be multi-dimensional 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 interim-distribution) but not on his identity. It is a reasonable assumption which is satisfied by the majority of social choice sets. For instance, in many mechanism design papers, a mechanism is efficient if it implements the utilitarian social choice set, which satisfies our definition of Anonymity³.

The Walrasian correspondence is another example of a well-known social choice set which satisfies Anonymity. The issues of the existence of equilibrium and its welfare properties in economies with adverse selection have been analyzed by many papers in

³See for example the papers by Mezzetti (2004), Jackson and Sonnenschein (2007).

the context of the Walrasian mechanism⁴. It has been shown that the equilibrium, if it exists, is inefficient. Since the usual justification for competitive behavior is the large number of agents in both sides of markets (indeed, most of these papers assume a continuum of agents), one can apply our mechanism in order to implement the full-information competitive equilibrium allocations in the examined economies.

Moreover, it should be pointed out that the assumption of the interim distribution of types being common knowledge is needed because we consider general social choice sets. 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 interim distribution⁵. Our mechanism can still truthfully implement the desirable allocations, given that the social planner knows the interim distribution. 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.

Finally, the issues of participation constraints and of ex-post feasibility (off-the-equilibrium-path feasibility) are also examined in the paper. The only point we can make at this point is that we address these concerns with slight modifications of the main mechanism. Ex-post participation (off-the-equilibrium-path participation), however, may indeed be violated by the mechanism we present in this paper, but this is an event that, in equilibrium, should happen with zero probability.

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 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

⁴Examples include 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).

⁵E.g. the Walrasian correspondence in the Rothschild-Stiglitz model.

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 set. 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 information 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 interim distribution of types (when the signal is perfect)⁶. 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 interim-distribution, 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 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

⁶In a sense, in our model agents receive private signals as well, but one can think of them as perfect signals about the interim distribution. As we have already mentioned, a small noise about the precision of these signals does not alter our results.

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 set of potential types. The vector θ 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. Φ is the ex-ante cumulative distribution function over the set of all possible type-profiles Θ , with $\Phi(\theta)$ the ex-ante probability that the type-profile θ is realized.

S is the set of all states. Each state s is a complete description of the world, including the economic characteristics of each agent. This means that the state describes agents' features, such as preferences, productivity, individual endowments or any other economically pertinent information. The probability distribution over states Π is a function of the type-profile θ . Therefore, $\pi(s|\theta)$ is the probability of state s arising, conditional on the type-profile θ .

β is an unordered collection of I realization of types (potentially the same types for some realizations). One interpretation is that β is the distribution of types that have been realized. Given a β , the exact number of agents who have a specific type is known for all types. We slightly abuse terminology by calling β the interim distribution of types. $\Theta(\beta)$ is the set of all type-profiles consistent with the interim distribution 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 $L \times S \geq 2$. L can be interpreted as the number of commodities in the economy. Each a is an S -tuple of feasible state-dependent allocations. Furthermore, we impose the following two restrictions on preferences. First, we assume that preferences are represented by expected utility functions:

$$U_i(a) = \sum_{\theta_{-i}} \left[\sum_{s \in S} u_i(a, s) \pi(s|\theta_i, \theta_{-i}) \right] \phi(\theta_i, \theta_{-i}), \quad \theta_{-i} \in \Theta_{-i}(\beta|\theta_i)$$

$U_i(a)$ is the expected utility to agent i when he receives allocation a , with $u_i(a, s)$ the decision-outcome payoff in state s (preferences may be state-dependent) and θ_{-i} is a

type-profile for all agents, excluding i , which is consistent with the interim-distribution of types β . Second, for all agents, preferences satisfy the local common-indifference property. This is a requirement that the intersection between 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 1: Let $C_{i\epsilon}(a) = \{c \in A : U_i(c|\theta_i, \theta_{-i}) = U_i(a|\theta_i, \theta_{-i}), \|c - a\| < \epsilon\}$. The **Local Incentive Compatibility Property** is satisfied if $\forall i \in I, \forall a \in A$ and $\forall j \in I, \theta_j \neq \theta_i$:

either i) $\exists \bar{\epsilon}_{ij} > 0 : \dim(C_{i\epsilon}(a) \cap C_{j\epsilon}(a)) \leq L \times S - 1, \forall \epsilon < \bar{\epsilon}_{ij}$

or ii) $\forall \epsilon > 0 : \dim(C_{i\epsilon}(a) \cap C_{j\epsilon}(a)) \leq L \times S - 1$

LICP 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 LICP but it violates the SCP. Also, LICP 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 LICP is also satisfied as part ii) of the definition implies⁷. Below we provide two diagrams, which illustrate the LICP and distinguish it from the SCP.




Figure 1: Indifference Curves satisfying LICP

⁷Note that we could alternatively characterize these restrictions 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 Incentive Compatibility. In this case, the only difference from the definition provided above is the definition of $C_{i\epsilon}(a)$: $C_{i\epsilon}(a) = \{c \in A : c \sim_i a, \|c - a\| < \epsilon\}$.

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.

4 Implementation of First Best Allocations

4.1 Full Implementation

This section presents the main result of the paper. We claim that if the interim-distribution of types is common knowledge and the social choice set satisfies Pareto efficiency and Anonymity, then a mechanism exists that fully implements it. First, we define our concept of Anonymity for Social Choice Sets and then we present a series of Lemmata, which are used in the proof. The proofs of the Lemmata are presented in the Appendix. Notice also that we eschew away from the issue of off-the-equilibrium-path feasibility for now. We deal with it in the subsequent subsection.

Definition 2: A Social Choice Set satisfies Anonymity if, for every social choice function in the set, each agent's assigned allocation depends on his type and the interim-distribution of types: $a_i^* = a(\theta_i, \beta)$.

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 interim-distribution of types there is a unique collection of allocations to be assigned to agents. The order of the allocations does depend on the type-profile θ , but the collection of individual allocations is the same for all type-profiles consistent with the same interim-distribution.

It is also noteworthy to mention 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 function⁸.

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: $U_i(a_i) \geq U_i(a_j), \forall j \in I$.

⁸In this paper, we prove that Anonymity and LIPC are sufficient conditions for first-best implementation, but we have little to say about necessity.

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: $U_j(a_j) \geq U_j(a_i), \forall j \in I$.

Proof: See the Appendix

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

Proof: See the Appendix

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 $\mathbf{Rank}(\mathbf{K}) = \{i \in I : U_i(a_i) \geq U_i(a_j), \forall j \in I\}$, 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 $\mathbf{Rank}(\mathbf{K}-1) = \{i \in I - \mathbf{Rank}(\mathbf{K}) : U_i(a_i) \geq U_i(a_j), \forall j \in I - \mathbf{Rank}(\mathbf{K})\}$. By iteration, we can define K groups, $1 \leq K \leq I$, 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 rank⁹. We will also refer to group $\mathbf{Rank}(\mathbf{K})$ as the group with the **highest rank** and group $\mathbf{Rank}(\mathbf{1})$ as the group with the **lowest rank**. We will shortly exploit this ranking of agents on Pareto efficient allocations in order to prove our main claim. Another result required for the proof comes from the LICP and it is provided in Lemma 3.

Lemma 3: If the LICP holds, then around 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.

⁹One extreme case is when an allocation exhibits no-envy, in which case $\mathbf{Rank}(\mathbf{K})$ 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

In effect, Lemma 3 states that 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. This property, 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 interim distribution, 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. We construct this argument formally in the proof of Proposition 1.

Proposition 1: Assume that the economy E satisfies the LICP and that the interim-distribution of types β is common knowledge. Then, for every allocation $a^* \in PF(E)$, which satisfies Anonymity, there exists a mechanism, for which a is the unique Bayes-Nash equilibrium allocation and agents truthfully report their private information.

Proof: The proof is done by construction. Let $a^*(\theta)$ be the first-best allocation which is to be implemented for each type-profile of types, with $a_i = a_i^*(\theta_i, \beta)$. Also, let $\lambda_\theta(\beta)$ and $\lambda_\theta(\mathbf{m})$ be the number of agents of type θ according to the interim distribution β and the received messages \mathbf{m} , respectively.

Each agent reports his type m_i and a final allocation is received according to the following:

- i) If $\mathbf{m} \in \Theta(\beta)$, then each agent receives $a_i^*(m_i, \mathbf{m}_{-i})$.
- ii) If \mathbf{m} is such that for two types, (θ, θ') , the number of reported agents is different from number of agents in the interim-distribution ($\lambda_\theta(\mathbf{m}) \neq \lambda_\theta(\beta)$, $\lambda_{\theta'}(\mathbf{m}) \neq \lambda_{\theta'}(\beta)$), then agents receive incentive compatible allocations around the first best allocation of the type with the lower rank between θ and θ' . If θ and θ' are of the same rank, then agents receive incentive compatible allocations around the first-best allocation of one of the two types arbitrarily (say the type with the lowest index θ).
- iii) If $\lambda_\theta(\mathbf{m}) \neq \lambda_\theta(\beta)$ for three types or more, then agents receive the incentive compatible allocations around the first-best allocation of a type with the lowest rank.

Under the mechanism above, it is a strictly dominant strategy for all agents of types of rank(K) to truthfully report their type. To see this consider the different beliefs of an agent of rank(K) (say i) about the messages that other agents will send. If i believes that all other agents will truthfully report their type then the best-response for him is clearly to truthfully report as well (he has nothing to gain by misrepresenting his preferences as he envies no other type's allocation).

If i believes that only one agent, of a different type, will misreport her preferences, then he still prefers to truthfully report his type, irrespectively of the rank of the other agent. Say that i believes that j is of the same rank as him but of different type and that she will misrepresent her preferences as being of type θ_i . If i reports that he is of type θ_j , then the two lies will cover each other and i will receive a_j^* . But if he chooses to report θ_i , then $\lambda_{\theta_j}(\mathbf{m}) \neq \lambda_{\theta_j}(\beta)$ and $\lambda_{\theta_i}(\mathbf{m}) \neq \lambda_{\theta_i}(\beta)$, in which case he receives an incentive compatible allocation around a_j^* . By construction then, i prefers to truthfully report his type. The same argument goes through if j is of lower rank than i .

Finally, in the case where i believes that multiple misrepresentations will take place, either in types of rank(K), or in other ranks, then, irrespectively of his message, $\mathbf{m} \neq \beta$ (if all representations but one cancel out then we go back to the analysis of the previous case). In this case, he still prefers to report his type truthfully in order to receive an incentive compatible allocation. Therefore, it is a strictly dominant strategy for an agent of rank(K) to truthfully report his type, irrespectively of his beliefs about what other agents will do.

Given this, then it is a best response for an agent of rank(K-1) to also truthfully report his type. Say that agent i , who is of rank(K-1), envies the allocation of some type θ_j of rank(K). Of course, if i believes that some agent of type θ_j will report as being of type θ_i , then the best response for i is $m_i = \theta_j$, but, as we showed, this cannot be an equilibrium¹⁰. Hence, if i believes that all agents will truthfully report, he prefers to truthfully report 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 truthfully reveal his type, 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 truthfully report, agents of rank(K-1) also truthfully report.

By induction, we conclude that for an agent of rank κ , if all agents of higher rank are expected to truthfully report 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 truthfully report, then the only possible equilibrium is when all agents truthfully report. Therefore, the unique Bayes-Nash equilibrium of the mechanism is for all agents to truthfully report their type and to receive the allocation $a_i^*(\theta_i, \beta), \forall i \in I$. ■

¹⁰This argument also makes clear that our paper is not one of dominant strategy implementation, as only rank(K) individuals have dominant strategies.

The result depends crucially on the fact that the rank of types is known. This is due to the interim-distribution 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 truthfully reveal his type. 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 LICP is also required for the uniqueness of the equilibrium, as it allows for agents to strictly improve their payoff if they truthfully report their type. Once again, if LICP is violated, then one can still easily construct mechanisms which implement the first-best allocations, but the uniqueness of the equilibrium may not be possible. Therefore, the common knowledge of the interim-distribution, Anonymity and LICP are jointly sufficient conditions for full implementation of first-best allocations, but they are not necessary.

4.2 Off-the-equilibrium-path feasibility

We now turn to the issue of the feasibility of the mechanism out of the equilibrium path. This is a valid concern since the mechanism requires that, if a lie is detected, all agents, who reported the “over-reported” type receive almost identical allocations. That is, if type θ is over-reported and type θ' is under-reported, all agents who reported types θ and θ' need to receive allocations close to $a_{\theta'}^*$. This may not be feasible.

In this subsection we show why feasibility is not violated in many cases of economic interest. We showcase this by providing examples of adverse selection economies, which have been widely used in the literature, and for which our mechanism is feasible. We also provide sufficient conditions for ensuring feasibility.

Before we proceed, we need to clarify two points. First, feasibility does not affect the ability of our mechanism to implement the desirable allocations but rather the uniqueness of the equilibrium. Or, to be more accurate, if we explicitly require the mechanism to be feasible off-the-equilibrium-path, then there may be multiple equilibria, some of which involve untruthful reporting. But the truthful equilibrium will always be one of them.

To see this, suppose that, in our original mechanism, it is infeasible for all agents of type θ and θ' to receive allocations close to $a_{\theta'}^*$. But the allocation $a^m(\theta, \theta')$, which gives the minimum quantity between $a_{\theta l_s}^*$ and $a_{\theta' l_s}^*$ for each state-contingent commodity l_s , is feasible by construction. If agents are to receive $a^m(\theta, \theta')$, whenever the numbers of reports for θ , θ' are different from the interim-distribution, then the mechanism is feasible off-the-equilibrium-path. Moreover, if an agent of type θ expects everyone else to truthfully report, it is a best-response for him to truthfully report as well, because

$U_i(a^*(\theta)) > U_i(a^m(\theta, \theta'))$.

If, however, $U_i(a^m(\theta, \theta')) > U_i(a^*(\theta'))$, then the mechanism has other equilibria, where i misreports his type as θ' if he believes that some agent of type θ' will report type θ instead. One can think of these undesirable equilibria as “coordination failures”. Agents of higher rank fear that, if they truthfully report (while others don’t), then they will be severely punished, if a lie is detected. They prefer to cover the lie, instead, in order to receive the allocation of the misreporting agent rather than the off-the-equilibrium punishment.

Second, the issue of feasibility can be easily overcome, without generating multiplicity of equilibria, if it concerns agents of the same rank. Consider the following modification of our mechanism. Suppose that types θ and θ' are of the same rank and that $\lambda_\theta(\mathbf{m}) \neq \lambda_\theta(\beta)$ and $\lambda_{\theta'}(\mathbf{m}) \neq \lambda_{\theta'}(\beta)$. Without loss of generality, suppose that $\lambda_\theta(\mathbf{m}) = m > \lambda_\theta(\beta)$ and $\lambda_{\theta'}(\mathbf{m}) = n < \lambda_{\theta'}(\beta)$, with $n + m = \lambda_\theta(\beta) + \lambda_{\theta'}(\beta)$. The modified mechanism specifies that the n agents, who report θ' , receive $a_{\theta'}^*$ and the m agents, who report θ , are arbitrarily (in terms of a stochastic process or a lottery) assigned to $\lambda_\theta(\beta)$ allocations a_θ^* or to $\lambda_{\theta'}(\beta) - n$ allocations $a_{\theta'}^*$. In the case where $\lambda_\theta(\mathbf{m}) = \lambda_\theta(\beta)$ and $\lambda_{\theta'}(\mathbf{m}) = \lambda_{\theta'}(\beta)$, then agents receive the allocation of the respective type they report.

This mechanism satisfies feasibility, as far as off-the-equilibrium allocations of agents of the same rank are concerned. This is because it shuffles $\lambda_\theta(\beta) + \lambda_{\theta'}(\beta)$ allocations (which are feasible by the definition of Pareto efficiency) to $\lambda_\theta(\beta) + \lambda_{\theta'}(\beta)$ agents, according to their reports. Furthermore, it remains a strictly dominant strategy for both types to report truthfully, irrespectively of what the other agents do. This is because, irrespectively of what other agents report, an agent of type θ (θ') maximizes his chances of receiving his most preferred allocation by reporting truthfully. The mechanism can also be easily extended to deal with multiple types of agents of the same rank. We make this point because we want to highlight that the feasibility concerns are most difficult to solve across ranks. Below we provide sufficient conditions for dealing with issues of cross-rank feasibility:

1. for every pair of agents $\{i, j\}$, with types θ and θ' and ranks κ and κ' respectively, $\kappa > \kappa'$, there exists some individual allocation $\hat{a}(\theta, \theta')$, such that $U_i(\hat{a}(\theta, \theta')) > U_i(a_{\theta'}^*)$ and $U_j(a_{\theta'}^*) > U_j(\hat{a}(\theta, \theta'))$, and $(\lambda_\theta(\beta) + \lambda_{\theta'}(\beta)) \times \hat{a}$ is feasible.
2. There exists some collection of Θ individual allocations $\{a_\theta^m\}$, such that $\forall \theta, \theta'$: $U_\theta(a_\theta^m) > U_\theta(a_{\theta'}^m)$, and, whenever $rank(\theta) > rank(\theta')$, $U_\theta(\hat{a}(\theta, \theta')) > U_\theta(a_{\theta'}^m)$. Furthermore, $I \times a_\theta^m$ is feasible for all θ .

The first condition ensures that higher ranks always prefer to truthfully report than to cover the lies of lower rank agents, given that all other agents truthfully report. The second condition ensures that if multiple lies are detected, then it is a best-response for one of the misreporting agents to deviate and report his true type. Both of these conditions are required for the uniqueness of the truthful equilibrium.

Note that, if LICP is satisfied, then the second condition is always satisfied by some allocation (think of the allocation $a^m\Theta$, which gives the minimum quantity between all $a_{\theta ls}^*$ for each state-contingent commodity ls). It is the first condition that may be violated under LICP alone.

In fact, there are many well known economic models, in which our mechanism can provide first-best solutions without violating the off-the-equilibrium-path feasibility. Spence (1973) and Rothschild-Stiglitz (1976) are two examples. Also, in cases where the number of agents of any particular type is relatively small compared to the total number of agents, feasibility is satisfied and the LICP is sufficient for full implementation. This is because, off-the-equilibrium-path, one can provide allocations very close but slightly less than the first-best allocations for those types, for which reports match the interim-distribution, and use the retained surplus to provide the off-the-equilibrium allocations for the remaining two types. This is relevant for models of adverse selection with a continuum of types, which we approximate with our model by letting Θ be arbitrarily large.

Finally, in some cases of interest, feasibility issues may not be a concern because of the existence of “external surplus”, which can be used to make the off-the-equilibrium allocations credible. For example, in the case of the store manager of the introduction, it is plausible to assume that, apart from the orders of his customers, he has other goods in stock, which he can provide in case there is greater demand for a certain bundle than the number of orders. And, even in the case where the available stock is not sufficient for these purposes, he can still order and deliver the goods at a later date. What matters from the perspective of the customers is that the off-the-equilibrium allocations are credible, even though they never materialize in equilibrium.

4.3 Robustness to small perturbations

So far we have assumed that the interim-distribution of types is 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 B be the set of all possible interim-distributions that can be generated by Θ . By assumption, $\bigcup_{\beta \in B} \Theta(\beta) = \Theta$. Suppose, now, that there is a small noise about the probability of the interim-distribution. Agents have a probability distribution over the set of interim-distributions. With probability $1 - \sum_{\gamma \in B} \epsilon_\gamma$, the interim-distribution β will be realized, while ϵ_γ is the probability that some other interim-distribution $\gamma \in B$ will be realized, with $\epsilon_\gamma > 0, \forall \gamma \in B$.

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 the interim distribution:

$$\begin{aligned}
U_i(a) &= (1 - \sum_{\gamma \in B} \epsilon_\gamma) \sum_{\theta_{-i} \in \Theta_{-i}(\beta|\theta_i)} \left[\sum_{s \in S} u_i(a, s) \pi(s|\theta_i, \theta_{-i}) \right] \phi(\theta_i, \theta_{-i}) \\
&\quad + \sum_{\gamma \in B} \epsilon_\gamma \left[\sum_{\theta_{-i} \in \Theta_{-i}(\gamma|\theta_i)} \left[\sum_{s \in S} u_i(a, s) \pi(s|\theta_i, \theta_{-i}) \right] \phi(\theta_i, \theta_{-i}) \right]
\end{aligned}$$

We also assume that for each $\gamma \in B$ and for every θ_i corresponds an individual allocation $a_i^*(\theta_i, \gamma)$ such that any I-collection of individual allocations is consistent with γ , Pareto optimal and satisfies Anonymity. In other words, for every γ there is a set of Pareto-optimal allocations to be implemented, each one corresponding to a specific realization of a type-profile θ consistent with γ and Anonymity.

In the case of uncertainty about the interim distribution, 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 one 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 lowest rank under β (and potentially other ranks for other γ 's). If he knows that β is the interim distribution with certainty, then under the mechanism presented in the previous sub-section, he would strictly prefer to report his type truthfully than report any other type:

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

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

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

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 truthfully report his type. The argument can be repeated for any other agent j of different rank according to β . Given a sufficiently small vector of probabilities ϵ , j expects all lower-rank agents to truthfully report and his best-response is to truthfully report as well, irrespectively of the messages send by agents of the same or higher rank. Hence, there exists some vector ϵ , 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 interim distribution of types is uncertain but there is a sufficiently high probability that some distribution β will be realized, then the mechanism of Proposition 1 fully implements the first-best allocations for every interim-distribution.

Proof: It follows from the analysis above.

It is noteworthy to point out 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 interim distribution γ . In other words, the almost certainty about β makes agents to report their type truthfully irrespectively of the interim distribution that is eventually realized. As a consequence, agents receive first-best allocations for all realized interim-distributions. 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 the interim distribution.

4.4 Convergence to ex-ante distributions

So far we have shown our main result and that it is robust to small uncertainty about the interim distribution. We also want to show that if the number of agents becomes very large then the interim-distribution converges to the ex-ante distribution of types, in which case our informational assumptions converge to the standard assumptions in the adverse selection literature, i.e. agents know the ex-ante probability of each type occurring. This allows us to relate our formulation and result to large economies with adverse selection problems, and make the claim that in this economies, because the interim-distribution is effectively common knowledge, one can implement first-best allocations. However, in order to make this point we need to reformulate certain aspects of E .

First, we have explicitly introduced the probability function Φ over the space of type-profiles, but we have said little about the unconditional probability of each type. In order to make our formulation comparable to the rest of the literature, in this subsection, we restrict Φ so that each type θ has the same unconditional probability of occurring even though types may not be independently distributed¹¹. Given this restriction, the ex-ante probability of a type θ is given by:

$$\tau(\theta) = \sum_{\theta_{-i}} \phi(\theta, \theta_{-i}), \theta_{-i} \in \Theta_{-i}$$

Furthermore, we need to impose restrictions on Φ for comparing economies with a different number of agents. Since each economy has a different number of type-profiles and interim-distribution for the same set of types, if the number of agents increases, we need to index the distribution of type-profiles by the number of agents in the economy: Φ_I . For consistency and comparability purposes, therefore, we impose that the unconditional (ex-ante) probability of a type in all economies is the same:

¹¹In fact, the formulation of section two is more general by allowing the probability of a type to depend on the identity of an agent so that some agents may face different probabilities of receiving a specific type than others. However, the restriction we are imposing in this section is purely for comparability purposes.

$$\sum_{\theta_{-i}} \phi_I(\theta, \theta_{-i}) = \sum_{\theta_{-i'}} \phi_{I'}(\theta, \theta_{-i'}) = \tau(\theta) , \forall I, I' \in \{1, 2, \dots, \infty\}$$

Then, given the above restrictions, one can apply the Weak Law of Large Numbers and conclude that as the number of agents in the economy becomes infinitely large, the relative frequency of occurrences of a certain type in the population converges to the ex-ante probability of the type:

$$\lim_{I \rightarrow \infty} \left(\frac{\lambda_{\theta}(\beta_I)}{I} \right) = \tau(\theta)$$

But this is exactly the information provided by the interim-distribution: the number of agents, for whom type θ has realized. Hence, at the limit, the relative frequency of types in the population (interim-distribution) must coincide with the ex-ante probability whenever the distribution of types does not depend on agents' identity (which is the assumption made in the literature).

4.5 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 analysis¹².

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 truthfully reveal their type and they receive a first-best allocation. Our result relies on information aggregation and appropriately chosen punishments. If the interim distribution is known (perfectly or imperfectly), then one can aggregate the messages that all agents are sending out and uncover any misreporting(s), even if the identity of the liar is not known.

¹²Of 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 allocations 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.

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 also be pointed out that the assumption of the interim distribution of types being common knowledge is needed because we consider general social choice sets. 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 interim distribution. Our mechanism can still truthfully implement the desirable allocations, given that the social planner knows the interim distribution. 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.

References

- Akerlof, G. (1970): “The Market for “Lemons”: Quality Uncertainty and the Market Mechanism”, *The Quarterly Journal of Economics*, 84, 488-500.
- 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): “Equilibria and Pareto Optima of Markets with Adverse Selection”, *Economic Theory*, 7, 207-235.
- Jackson, M. (1991): “Bayesian Implementation”, *Econometrica*, 59, 461-477.
- Jackson, M. and H. Sonnenschein (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 Weltwirtschaft an der Universität Kiel.
- 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.
- Palfrey, T., and S. Srivastava (1989): “Implementation with Incomplete Information in Exchange Economies”, *Econometrica*, 57, 115-134.
- 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: $U_i(a_i) \geq U_i(a_j), \forall j \in I$.

Proof: Take any allocation $a \in PF(E)$ and suppose that the claim of Lemma 1 does not hold. Then, this means that there is no agent, who does not envy the allocation of some other agent. Therefore, for every $i \in I$, there exists some $j \in I$, with $j \neq i$, such that $U_i(a_j) > U_i(a_i)$. Since this holds for all agents in I , then there must be some reassignment of allocations, such that a subset of I receives the allocations that they envy (while the rest of the agents retain their original allocation) and hence some subset of I can be made better-off, while the rest remain as well-off as under a .

To see this, consider the following algorithm. Take arbitrarily one agent i and reassign to him one of the allocations that he envies, say a_j . Then move on to the agent j , from whom the allocation has been reassigned, and check the set of the allocations that he envies as well. If a_i is contained in the set of allocations envied by j , then reassign that allocation to j and stop searching for reassignments. Otherwise, arbitrarily reassign to him some allocation a_h and move to agent h . Repeat the same procedure until the set of allocations envied by some agent k in this sequence contains the allocation from some agent l , who has already being reassigned some other allocation and reassign a_l to agent k . If $l = i$ then stop the allocation reassignment. If $l \neq i$, then give to all agents from i until l (the agents who were at the beginning of this sequence of reallocations until reaching agent l) their original allocation and leave the rest of the reassignments unchanged. This is done in order to ensure that no allocation is to be reassigned to more than one agent, which would violate feasibility.

Notice that, because, by assumption, all agents envy the allocation of at least one other agent, the above algorithm can at most generate a sequence of I reassignments. Since all reassignments are made to agents within the same set I , then at some point the reassignment will lead to an allocation of an agent which has already being reassigned. This implies that there is a re-arrangement of the I allocations, such that some agents in I are made strictly better-off than in the original allocation. But this means that a is not Pareto-optimal, which is a contradiction. ■

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: $U_j(a_j) \geq U_j(a_i), \forall j \in I$.

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 been 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 $\widehat{I} \subseteq I$ and let $\widehat{A} = \{a_i : i \in \widehat{I}\}$. Then, if $a \in PF(E)$, Lemma 1 and 2 hold for \widehat{I} with regards to \widehat{A} as well.

Proof: Take any subset of agents \widehat{I} of the set I. Suppose that Lemma 1 and 2 do not hold over the set \widehat{A} , which is the set of individual allocations of the agents in \widehat{I} . Then, it is possible to find a reassignment of allocations between the agents in \widehat{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: If the LICP holds, then around 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.

Proof: Recall that $C_{i\epsilon}(a) = \{c \in A : U_i(c|\theta_i, \theta_{-i}) = U_i(a|\theta_i, \theta_{-i}), \|c - a\| < \epsilon\}$. Also, define $L_j(a_i)$ to be the lower-contour set of agent j associated with allocation a_i : $L_j(a_i) = \{c \in A : U_j(c|\theta_j, \theta_{-j}) < U_j(a_i|\theta_j, \theta_{-j})\}$ and $V_j(a_i)$ to be the upper-contour set: $V_j(a_i) = \{c \in A : U_j(c|\theta_j, \theta_{-j}) > U_j(a_i|\theta_j, \theta_{-j})\}$.

\mathbf{H} is a $L \times S - 1$ hyper-plane, which passes through a_i , and is perpendicular to the MRS of some type's indifference curve, which also passes through a_i . \mathbf{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 LICP, there exists some $\bar{\epsilon} > 0$ such that for every $\epsilon < \bar{\epsilon}$, 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 picture below).




Figure 2: LICP and Local Incentive Compatibility

Say that agent k is the type with the smallest upper contour set within ball $B_\epsilon(a_i)$ and subspace A_1 : $V_k(a_i) \cap B_\epsilon(a_i) \cap A_1 \subset V_l(a_i) \cap B_\epsilon(a_i) \cap A_1, \forall l \in \Theta$. Then, 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. ■