Available online at www.sciencedirect.com
ScienceDirect
Journal of Economic Theory 136 (2007) 476–496
JOURNAL OF Economic Theory
www.elsevier.com/locate/jet

Mechanism design with multidimensional, continuous types and interdependent valuations

Nolan H. Miller[a,∗,1], John W. Pratt[b], Richard J. Zeckhauser[a], Scott Johnson[c,†]

[a] *Kennedy School of Government, Harvard University, 79JFK St, Cambridge, MA 02138, USA*
[b] *Harvard Business School, Harvard University, Boston, MA 02163, USA*
[c] *Australian National University, Canberra ACT 0200, Australia*

Received 6 December 2004; final version received 11 September 2006
Available online 7 November 2006

Abstract

We consider the mechanism design problem when agents' types are multidimensional and continuous, and their valuations are interdependent. If there are at least three agents whose types satisfy a weak correlation condition, then for any decision rule and any $\varepsilon > 0$ there exist balanced transfers that render truthful revelation a Bayesian ε-equilibrium. A slightly stronger correlation condition ensures that there exist balanced transfers that induce a Bayesian Nash equilibrium in which agents' strategies are nearly truthful.

JEL classification: C70; D60; D70; D82

Keywords: Mechanism design; Interdependent valuations; Multidimensional types

1. Introduction

Eliciting private information to guide social decisions is a classic problem of economic theory. For the private-values case, the pioneering work of Vickrey [32], Clarke [3], and Groves [14] shows that if each agent's preferences depend only on his own information and if the budget need not balance, externality payments make honest revelation a dominant strategy. However, Green and Laffont [12,13] show that dominant strategy implementation is generally incompatible

∗ Corresponding author. Fax: +1 617 496 5747.
E-mail address: nolan_miller@harvard.edu (N.H. Miller).

1 A previous version of some results in this paper circulated under the title "Efficient design with multidimensional, continuous types and interdependent valuations."

† Deceased.

0022-0531/$ - see front matter
doi:10.1016/j.jet.2006.09.005

with the requirement that the transfers balance the budget. If the solution concept is weakened, positive results are possible. For example, d'Aspremont and Gérard-Varet [7,8] show in the private-values environment that if agents' beliefs about other agents' types satisfy a certain condition, which they call compatibility, then for any efficient decision rule there exist balanced Bayesian incentive-compatible transfers that implement it. [2] Later, d'Aspremont et al. [10] show that when there are three or more agents the compatibility condition is generically true, and hence that for generic distributions of agents' types there exists a Bayesian incentive-compatible, Pareto-optimal mechanism.

The mechanism design problem has proved more challenging in the case of interdependent valuations, i.e., when one agent's private information affects other agents' preferences. Dasgupta and Maskin [11] study auctions with interdependent valuations and show that a generalized Vickrey auction is efficient if bidders' types are one dimensional and satisfy a single-crossing property. In general mechanism-design problems, positive results have been mostly limited to the case where agents' types take on only finitely many values. Work in this area includes Johnson et al. [19]; Matsushima [23,24]; and McLean and Postlewaite [26]. Crémer and McLean [5,6] study the related question of when it is possible for the designer to earn as much profit as he would were he able to observe the agents' realized private information—the so-called full surplus extraction problem—and show that full extraction is possible when agents' types are suitably correlated. Aoyagi [1] considers a model with finite type sets and interdependent valuations and shows that if the distribution of agents' types satisfies a dependence condition similar to ours, then for any decision rule there exists a balanced, Bayesian incentive-compatible mechanism that implements it.

When types are multidimensional and continuous and valuations are interdependent, the problem becomes even more difficult. After their possibility result for the one-dimensional case, Dasgupta and Maskin [11] go on to show that when bidders' types are multidimensional and independently distributed there may be no efficient auction. In a general mechanism design framework, Jehiel and Moldovanu [18], henceforth JM, explore the difficulties of Bayesian incentive-compatible (BIC) implementation of efficient decision rules when types are multidimensional and continuous and valuations are interdependent. They show that when agents' types are independently distributed, efficient BIC design is possible only when a certain "congruence condition relating the social and private rates of information substitution is satisfied" [18, p. 1237]. In effect, this congruence condition requires that there be one agent whose relative preference over any two alternatives remains constant for all values of that agent's information that make the social planner indifferent between those alternatives. They then show that when types are multidimensional the set of payoff functions that satisfy this condition is non-generic, implying that efficient BIC design is generally impossible. [3]

The present paper addresses the mechanism design problem in environments in which agents' private information is continuous, multidimensional, and mutually payoff-relevant (i.e., valuations are interdependent). However, we relax the JM assumption that agents' private information is independently distributed. Our primary interest is to show that when there are three or more agents and agents' types are stochastically dependent it is possible to design a system of

[2] Unlike its use in implementation theory (see [16]), throughout the paper we use "implement" to refer to the case where there is an outcome of the game that agrees with the decision rule.

[3] Although Jehiel and Moldovanu [18] focuses on the impossibility of efficient BIC design, much of the importance of the result lies in the fact that it implies that robust mechanism design using belief-free concepts such as ex post equilibrium is also impossible. We return to this point in Section 5.

budget-balanced transfer payments that induces agents to (nearly) truthfully reveal their private information and that (nearly) implements any decision rule. In our first result (Theorem 1), we show that under a mild dependence condition on the distribution of agents' types, which we call Stochastic Relevance, for any $\varepsilon > 0$ there exist budget-balanced transfers such that truthful revelation is an ε-best response to other agent's truthful announcements. Thus the strategy profile where all agents announce truthfully is a Bayesian ε-equilibrium. In our second result (Theorem 2), we show that a slightly stronger version of Stochastic Relevance, which we call Uniform Stochastic Relevance, ensures that for any $\delta > 0$ there exist balanced transfers under which there is a Bayesian Nash equilibrium (BNE) of the announcement game in which the distance between agents' equilibrium announcements and their true types is no more than δ, i.e., that there is a nearly truthful BNE. Thus our results provide a complement to those of JM. When the distribution of agents' types satisfies our dependence assumptions, then incentive-compatible design is possible. Further, our implementation results place very few additional requirements on agents' preferences. [4] In particular, we do not require a single-crossing property.

Mezzetti [28] considers implementation of efficient decision rules in a model in which the social planner bases transfers on agents' reports of both their types and the utility they realize from the social decision. The paper shows that implementation of efficient decision rules is generally possible using a two-stage Groves mechanism. However, since agents may not realize the utility from a social decision until long after the decision is made, this framework presupposes, among other things, that the planner is able to make long-term commitments to make transfers in the future. Even in circumstances where the two-stage mechanism is feasible, Mezzetti's results apply only to efficient decision rules, whereas the results of this paper apply to all decision rules. Further, the present paper imposes a more stringent form of budget balance than Mezzetti. [5]

McAfee and Reny [25] (henceforth MR) consider the full surplus extraction problem in the case of continuous, multidimensional, and mutually payoff-relevant types with stochastically dependent information. Taking the game played by the agents as given, they show that it is possible to construct for each agent a finite menu of participation fee schedules that extracts almost all of the agent's rent from playing the game. However, they do not directly address the issue of which decision rules can be implemented, the primary concern of this paper. For example, with multidimensional types and interdependent values, there is, in general, no ex post efficient auction mechanism unless additional assumptions are made that ensure that the agents' multidimensional information can be summarized by a one-dimensional type [11,20,22]. Therefore, in such environments, the MR mechanism cannot extract the full information rent (i.e., the rents that would be generated if the auctioneer knew the agents' types), since the MR construction depends on the existence of an ex post efficient mechanism to which the participation fees can be appended. The present paper fills this gap by showing how to construct an ex post efficient mechanism in this environment. Appending the MR mechanism to ours then makes it possible to fully extract the agents' surplus.

Positive results on incentive-compatible implementation as well as full surplus extraction [1,5,6,25] rely on constructing a menu of lotteries for each agent such that the agent maximizes his expected utility when he chooses the lottery intended for his type. Intuitively, this is possible

[4] Specifically, we require only that agents' direct returns from the center's decision be bounded and suitably smooth.

[5] We adopt the standard definition in the literature that balanced transfers must sum to zero for any possible choice of actions by the agents. Mezzetti's transfers satisfy the weaker requirement that the transfers sum to zero on the equilibrium path when all players play truthful strategies.

whenever learning an agent's type provides information about the distribution of the other agents' types. Our analysis follows in the same spirit.

We capitalize on the literature in statistical decision theory on strictly proper scoring rules, which considers how an informed expert can be induced to truthfully reveal his beliefs about the distribution of future random events. [6] A scoring rule assigns payoffs to the expert based on his announced probabilities for various future events and the event that actually occurs. A strictly proper scoring rule has the property that the decision maker maximizes his expected score when he truthfully announces his beliefs about the distribution. Our incentive-compatibility results rely on using payments based on a proper scoring rule to drive agents toward truthful revelation of their private information. [7]

The paper proceeds as follows. Section 2 presents the model. Section 3 constructs scoring-rule payments that render truthful reporting a Bayesian ε-equilibrium. The basic construction is adapted in Section 4 to show that under a slightly stronger correlation condition similarly constructed payments ensure that there is an exact BNE in which agents' strategies are arbitrarily close to truthful. Section 5 discusses limitations of the approach in the paper, and Section 6 concludes. All proofs are presented in the Appendix.

2. The model

Suppose $I \geqslant 3$ agents, indexed by $i = 1, \ldots, I$, interact with the mechanism designer, whom we will call the center. The center's task is to elicit agents' private information in order to choose an alternative g from a set of alternatives G.

Each agent i has private information or type $t_i \in T_i$. Agent i's type space T_i is a non-empty, compact, convex subset of d^i-dimensional Euclidean space. For each i, d^i is a positive integer, and d^i may be different for different agents. We use $T = \times T_i$ to denote the product space of the I agents' type spaces. Following the standard notation we use $t = (t_1, \ldots, t_I)$ for the vector of types, t_{-i} for the vector of all but agent i's type, and t_{-ij} for all but the types of agents i and j.

Each agent's utility is quasilinear in his direct return from the social alternative, g, and money, x, taking the form: $u_i(t, g, x) = V_i(t, g) + x$. Note that agent i's direct return from alternative g depends on all agents' types. Hence valuations are interdependent.

A decision rule $g : T \to G$ maps a type for each agent to a social alternative. For simplicity, we assume that $g(t)$ is single valued. For $g(t)$ that is not single-valued, our implementation result applies to any selection from $g(t)$, and therefore this restriction is without loss of generality. Although we will impose a degree of smoothness on $g(t)$, we will not otherwise restrict it. In particular, we do not require that $g(t)$ be efficient.

We consider direct mechanisms in which each agent sends a message (announcement) to the center consisting of an element from his type space. We denote these announcements by $a_i \in T_i$, and let a, a_{-i}, and a_{-ij} refer respectively to the full announcement vector, the announcement vector leaving off agent i, and the announcement vector leaving off agents i and j. The remainder of the mechanism consists of a transfer function $x_i(a)$ for each i and a decision rule $g(a)$, with the standard interpretation that if the agents announce a, social alternative $g(a)$ is realized and transfer $x_i(a)$ is made to agent i.

An *announcement strategy* for agent i is a function $s_i(\cdot) : T_i \to T_i$ that specifies agent i's announcement in the message game as a function of his information. We will use the notation

[6] See Cooke [4] and the references therein for a discussion of scoring rules and their uses.

[7] Johnson et al. [19] employs a similar technique in the case of finite type and action spaces.

$s_i(\cdot)$ to refer to a strategy for agent i and $s_i(t_i)$ to denote to the announcement agent i makes under strategy $s_i(\cdot)$ when his type is t_i. Thus, $s_i(\cdot)$ is an element of a function space, while $s_i(t_i)$ resides in d_i-dimensional Euclidean space. We will use τ_i to denote the identity function on T_i, i.e., agent i's truthful strategy.

Denote the vector of transfer rules to all agents by $x(a)$, which we call a *transfer scheme*. A transfer scheme is *balanced* if its transfers sum to zero for all possible announcement vectors: $\sum_{i=1}^{I} x_i(a) = 0$ for all a. If a decision rule is implemented by a balanced transfer scheme, it requires no outside subsidy.

Since our mechanism is essentially the same for any decision rule and depends on the decision rule only through the direct return function, we integrate the decision rule into the direct return function, and then write $V_i(t, g(a))$ as $v_i(t, a)$. If there exists a transfer scheme that satisfies a particular solution concept with payoffs $v_i(t, a)$, then those transfers implement $g(a)$ under that solution concept. We make the following assumption regarding $v_i(t, a)$:

Assumption 1 (*Smooth direct returns*). For each i, expected direct returns are (jointly) twice continuously differentiable in a_i and t_i.[8]

Assumption 1 is not innocuous since it implies restrictions on the continuity of the underlying decision rule, $g(a)$, and on the set of possible decisions, G. Nevertheless, continuity seems to be a reasonable restriction in any situation that is appropriately modeled using continuous types. Further, discontinuous decision rules can often be approximated by continuous ones, and the results below would generalize to the case of decision rules that can be approximated by continuous rules.

Since $v_i(t, a)$ is continuous and T is compact, Assumption 1 implies that direct returns are bounded. Let $\bar{M} \geqslant 0$ denote the bound. That is, for any i, t_i and a, $|E\{v_i(t, a)|t_i\}| \leqslant \bar{M}$.

Types are distributed according to commonly known prior distribution $F(t)$ with support T. Let $f(t_j|t_i)$ be the density of agent j's private information conditional on agent i's private information, t_i, and let $f(t_{-i}|t_i)$ be the density of all other agents' private information conditional on agent i's private information. We impose two assumptions on agents' beliefs, a smoothness condition and a correlation condition.

Assumption 2 (*Smooth conditional distributions*). For each i and $j \neq i$, conditional densities $f(t_j|t_i)$ are jointly continuous in t_j and t_i and continuously differentiable in t_i. Similarly, $f(t_{-i}|t_i)$ are jointly continuous in t_{-i} and t_i and continuously differentiable in t_i.

We assume that the agents' private information is not independently distributed, which departs from the JM model. Specifically, our informativeness assumption, which we call *Stochastic Relevance*, is that the conditional distribution of the center's information be different for different values of each agent's private information.

Assumption 3 (*Stochastic relevance*). For each i, there exists an agent $j \neq i$ such that for any distinct types t_i and t_i' there exists $t_j \in T_j$ such that:

$$f(t_j|t_i) \neq f(t_j|t_i').$$

[8] In order to ensure the function's derivatives exist on the boundary of the domain, we assume $v_i(a, t)$ is defined and twice continuously differentiable on an open set containing T^2. An alternative approach would be to apply a unique extension theorem such as Proposition 7.5.11 in [30, p. 149]. We adopt the same approach in Assumption 2.

Let $\|\cdot\|_R$ denote the Euclidean norm, $\|t\|_R = \left(\sum_k (t_{ik})^2\right)^{1/2}$, where t_{ik} denotes the kth component of t_i, and let $\|\cdot\|_2$ denote the L_2 norm, $\|f\|_2 = \left(\int |f|^2 \, ds\right)^{1/2}$. We will write $f_j(\cdot|t_i)$ when we wish to denote agent i's beliefs about the distribution of t_j considered as a function. Lemma 1 follows as an immediate consequence of Assumptions 2 and 3.

Lemma 1. *Assumptions* 2 *and* 3 *imply that for each* i, *for any* $\delta > 0$ *there exists* $\mu > 0$ *such that*:

$$\left\|t_i - t_i'\right\|_R \geqslant \delta \quad \textit{implies} \quad \left\|f_j(\cdot|t_i) - f_j\left(\cdot|t_i'\right)\right\|_2 \geqslant \mu.$$

Taken together, Assumptions 2 and 3 and Lemma 1 imply that $f\left(t_j|t_i\right)$ and $f\left(t_j|t_i'\right)$ differ on an open subset of T_j and that $f_j(\cdot|t_i)$ and $f_j\left(\cdot|t_i'\right)$ are close together (as functions in L_2) if and only if t_i is close to t_i'. Thus they capture the idea that types should have similar beliefs if and only if they are close together. [9]

3. Existence of nearly Bayesian incentive compatible transfers

We begin by considering the question of whether there exist transfers that make the truth nearly a best response, provided that all other agents announce truthfully. Considering this question allows us to illustrate our construction in the simplest setting. In the next section, we go on to show that a similarly constructed payments establish that there is a nearly truthful BNE of the game.

We begin with the notion of ε-Bayesian Incentive Compatibility. [10] Transfer scheme $x(a)$ is *ε-Bayesian Incentive Compatible* (ε-BIC) if for any i, t_i, and a_i:

$$E\left\{v_i(t, t_i, t_{-i}) + x_i(t_i, t_{-i})\,|t_i\right\} \geqslant E\left\{v_i(t, a_i, t_{-i}) + x_i(a_i, t_{-i})\,|t_i\right\} - \varepsilon.$$

That is, if for each agent i, announcing truthfully is an *ε-best response* to the other agents' truthful announcements.

As discussed earlier, the mechanism we propose draws on the decision-theoretic literature on proper scoring rules. In particular, we employ the quadratic scoring rule. Suppose that agent j is using the truthful announcement strategy, $s_j(\cdot) = \tau_j$, and player i is being scored based on how well he predicts agent j's announced type. The quadratic score assigned to type t_j when agent i announces a_i is given by

$$Q\left(t_j|a_i\right) = 2f\left(t_j|a_i\right) - \int_{T_j} f\left(t_j|a_i\right)^2 \, dt_j.$$

Lemmas 2 and 3 establish basic properties of the quadratic scoring rule that will be used in the subsequent analysis.

Lemma 2. *For any agent* i, *choose an agent* j *according to Assumption* 3, *and suppose agent* j *truthfully announces his type*, $s_j(\cdot) = \tau_j$. *Truthful revelation uniquely maximizes agent* i*'s*

---

[9] Although it would add significant notational burden, Stochastic Relevance could be relaxed to allow for the case where agent i's beliefs about the joint distribution of a group of agents' types depends on t_i even though the marginal distribution for any other agent's type does not. Aoyagi [1, Assumption 2] presents such a condition for the finite case.

[10] ε-Bayesian incentive compatibility appears, for example, in [8].

expected quadratic score:

$$t_i = \arg\max_{a_i \in T_i} \int_{T_j} Q\left(t_j | a_i\right) f\left(t_j | t_i\right) dt_j.$$

As Selten [31] notes, the proof that truthful revelation uniquely maximizes the expected quadratic score also shows that the expected loss from agent i's announcing $a_i \neq t_i$ instead of his true type t_i is equal to the square of the L_2-distance between agent i's beliefs when his type is a_i and when his type is t_i. Lemma 3 exploits this property.

Lemma 3. *For any agent* i, *choose an agent* j *according to Assumption* 3. *For any* $\delta > 0$ *there exists* $\varepsilon > 0$ *such that the expected quadratic score for the distribution of agent* j*'s type from announcing* $a_i \neq t_i$ *with* $\|a_i - t_i\|_R \geqslant \delta$ *is at least* ε *worse than announcing truthfully:*

$$\|a_i - t_i\|_R \geqslant \delta \quad \textit{implies} \int_{T_j} Q\left(t_j | a_i\right) f\left(t_j | t_i\right) dt_j - \int_{T_j} Q\left(t_j | a_i\right) f\left(t_j | a_i\right) dt_j \geqslant \varepsilon.$$

Lemma 2 establishes that if the agents care only about the transfer, truthful announcement is agent i's unique best response when other agents' tell the truth. Lemma 3 ensures that there is no sequence of announcements far away from the truth whose expected scores converge to the expected score of the truth. This is needed in order to establish a uniform lower bound on the loss from an announcement that is far from truthful.

Our first main result shows that there exist ε-BIC, balanced transfers. The intuition is that in choosing whether to announce his true type or some other type the agent weighs the effects of lying on the expected transfer and on the expected direct return. If transfers are based on the quadratic scoring rule, then telling the truth maximizes the agent's expected transfer. However, since announcing truthfully does not necessarily maximize the expected direct return, the agent may have an incentive to deviate from truth-telling, sacrificing expected transfer in order to enjoy a personally superior social alternative. Of course, the agent's willingness to do so depends on how quickly the expected transfer declines relative to the increase in expected direct return. By scaling up the scoring-rule based payments to the agent, the center can increase the importance of the transfer loss relative to the direct return gain, making anything but a small deviation from the truth unprofitable.

Theorem 1. *Under Assumptions* 1–3, *for any decision rule and any* $\varepsilon > 0$ *there exist* ε-BIC, *balanced transfers*.

The essence of the proof is to divide agent i's announcements into two groups—those that are within δ of the truth and those that are not. Under the quadratic scoring rule, the expected transfer is maximized by telling the truth. Thus announcements that are within δ of the truth yield a smaller expected transfer but a possibly larger direct return. However, by choosing δ sufficiently small we ensure that the direct return gain from any announcement within δ of the truth must be less than ε. On the other hand, Assumptions 2 and 3 ensure that the loss in expected transfer from moving from a truthful announcement to one that is more than δ from the truth must be uniformly bounded away from zero, and thus scaling up the transfers increases the minimum loss in transfer from an announcement at least δ from the truth. Since direct returns are bounded, a sufficient scaling of the transfers ensures that the gain in direct return gain cannot outweigh the transfer loss, and thus that announcements that are at least δ from the truth must involve a total expected utility loss of at least ε.

The transfers are balanced using a permutation construction. That is, if agent 1 is given incentives to report truthfully by comparing his announcement to that of agent 2, then the transfer to agent 1 can be funded by a third agent (e.g., agent 3) without affecting any agent's incentive to report truthfully. Repeating this process for all agents balances the budget. Thus, while three or more agents are needed in order to balance the budget, if budget balance is not a concern, ε-BIC transfers exist with only two agents.

4. Existence of a nearly truthful Bayesian Nash equilibrium

Theorem 1 establishes that compensating agents using a sufficiently large scaling of the quadratic scoring rule renders truthful revelation an ε-best response, provided that the other agents announce truthfully. Although this idea has some intuitive appeal and makes the role of the quadratic scoring rule transparent, requiring agents to play merely ε-best responses rather than exact best responses begs the question of whether this limited rationality is necessary or merely a convenience. To address this concern, we next argue that, under reasonable conditions, payments based on a scaling of the quadratic scoring rule can be used to induce a BNE in which agents' strategies are arbitrarily close to the truth.

For a fixed transfer scheme $x(a)$, a BNE of the announcement game is a vector of strategies $(s_1(\cdot), \ldots, s_I(\cdot))$ such that for each i and t_i:

$$s_i(t_i) \in \arg\max_{a_i} E_{t_{-i}} \{v_i(t, a_i, s_{-i}(t_{-i})) + x_i(a_i, s_{-i}(t_{-i})) \,|t_i\}.$$

We endow the space of announcement strategies with the sup norm:

$$\left\| s_i(\cdot) - \hat{s}_i(\cdot) \right\|_{\sup} = \sup_{t_i} \left(\sum_{n=1}^{d_i} \left(s_{in}(t_i) - \hat{s}_{in}(t_i) \right)^2 \right)^{1/2}.$$

For $\delta > 0$, we call an announcement strategy, $s_i(\cdot)$, *δ-truthful* if $\|s_i(\cdot) - \tau_i\|_{\sup} \leqslant \delta$. That is, a δ-truthful announcement strategy is one in which the agent's announcement is always within distance δ of his true type. We say that a transfer scheme *δ-implements a decision rule in BNE* if under those transfers there exists a BNE in which all agents' strategies are δ-truthful. [11] Note that the concept of δ-implementation in BNE allows for the existence of BNE that are not δ-truthful.

Let C_i denote the space of continuous announcement strategies for agent i. For $\delta > 0$, let $C_i(\delta)$ be the space of continuous, δ-truthful announcement strategies for agent i: $C_i(\delta) \equiv \left\{ s_i(\cdot) \in C_i : \|s_i(\cdot) - \tau_i\|_{\sup} \leqslant \delta \right\}$. In the usual way, let $C(\delta)$ denote the product space $\times_i C_i(\delta)$, and let $C_{-i}(\delta)$ denote the product space of $C_j(\delta)$ for all agents except i, each endowed with the appropriate product topology.

The key step in constructing a δ-truthful BNE is ensuring that a version of Stochastic Relevance remains true even when agents' announcements are only δ-truthful. In order to ensure this we strengthen stochastic relevance as follows:

Assumption 4 (*Uniform stochastic relevance*). There exists $\phi > 0$ such that for each i, there exists an agent $j \neq i$ such that for any distinct types t_i and t_i' there exists an open ball $\theta_j\left(t_i, t_i'\right) \subset T_j$ with radius ϕ such that $f\left(t_j|t_i\right) \neq f\left(t_j|t_i'\right)$ for all $t_j \in \theta_j\left(t_i, t_i'\right)$.

[11] We may, on occasion, refer to single announcements as δ-truthful if for a particular t_i, $\|s_i(t_i) - t_i\|_R \leqslant \delta$ or to strategy profiles as being δ-truthful if each individual strategy is δ-truthful.

Stochastic Relevance (Assumption 3) implies that, for any distinct types t_i and t_i', $f\left(t_j|t_i\right)$ and $f\left(t_j|t_i'\right)$ differ on an open set of types for agent j. Uniform Stochastic Relevance (Assumption 4) strengthens Stochastic Relevance by requiring that there be a lower bound on the size of the open set on which $f\left(t_j|t_i\right)$ and $f\left(t_j|t_i'\right)$ differ that is independent of the particular pair of types t_i and t_i' that is chosen. It is straightforward to show that, by virtue of compactness, Assumption 3 implies the existence of such a uniform bound provided that t_i and t_i' are bounded away from each other, i.e., as long as there exists $\delta > 0$ such that $\left\|t_i - t_i'\right\|_R \geqslant \delta$. Thus, to the extent that Uniform Stochastic Relevance is stronger than Stochastic Relevance, it only restricts the behavior of beliefs as types t_i and t_i' become (arbitrarily) close together.

Since agent i's beliefs are continuous in t_i, as t_i and t_i' become very close, the two types' beliefs must also become very close. Uniform Stochastic Relevance rules out the case in which as t_i' converges to t_i the Lebesgue measure of the set of t_j where their associated beliefs differ, $\left\{t_j \in T_j | f\left(t_j|t_i\right) \neq f\left(t_j|t_i'\right)\right\}$, converges to zero. In other words, under Uniform Stochastic Relevance it cannot be that as t_i' approaches t_i, $f\left(t_j|t_i'\right)$ approaches $f\left(t_j|t_i\right)$ by becoming equal to it on an ever-larger set of t_j. Seen in this way, it is clear that many of the families of beliefs economists typically consider satisfy Uniform Stochastic Relevance. For example, beliefs where t_j is distributed normally with mean t_i satisfy Uniform Stochastic Relevance.

For an example of a family of beliefs that does not satisfy Uniform Stochastic Relevance, consider $T_i = T_j = [0, 1]^2$. Suppose that $f\left(t_j|t_i\right)$ is uniformly distributed on a disk centered at $t_j = t_i$ and having radius $1/10$ (for t_i suitably distant from the boundary of T_j). Consider $t_i = (1/2, 1/2)$. Let $\lambda(\cdot)$ denote Lebesgue measure. Since

$$\lim_{\left\|t_i' - (1/2,1/2)\right\|_R \longrightarrow 0} \lambda\left(\left\{t_j \in T_j | f\left(t_j|t_i\right) \neq f\left(t_j|t_i'\right)\right\}\right) = 0,$$

these beliefs violate Uniform Stochastic Relevance. This is because $f\left(t_j|t_i\right)$ and $f\left(t_j|t_i'\right)$ are equal on their common support, and as t_i' converges to t_i, the supports of $f\left(t_j|t_i\right)$ and $f\left(t_j|t_i'\right)$ converge as well. [12]

The existence of a uniform lower bound on how often the beliefs of two different types of agent i differ is important when agents' strategies are permitted to be δ-truthful (as in Theorem 2) rather than exactly truthful (as in Theorem 1). If $f\left(t_j|t_i\right)$ and $f\left(t_j|t_i'\right)$ are equal except for a very small set of t_j, then it is possible that, even though $f\left(t_j|t_i\right)$ and $f\left(t_j|t_i'\right)$ differ, the distribution of agent j's announcements resulting from a particular δ-truthful strategy (e.g., the δ-truthful strategy that is constant over the set of t_j where $f\left(t_j|t_i\right)$ and $f\left(t_j|t_i'\right)$ differ) is the same for t_i and t_i'. Lemma 4 shows that Uniform Stochastic Relevance ensures that different types t_i and t_i' have different beliefs about the distribution over a set of discrete events comprised of groups of announcements for some agent $j \neq i$, and that this difference remains even if agent j distorts his announcement slightly. [13]

[12] On the other hand, if $f\left(t_j|t_i\right)$ is distributed as a cone with a circular base of radius $1/10$ and peak at t_i, these beliefs would satisfy Uniform Stochastic Relevance since the set of points where the densities of $f\left(t_j|t_i\right)$ and $f\left(t_j|t_i'\right)$ are equal remains small (i.e., has Lebesgue measure zero) even as t_i and t_i' become arbitrarily close together.

[13] Lemma 4 is also useful for a more technical reason. When $s_j(\cdot) \in C_j(\delta)$, player j's announcement can be constant over an open interval. Hence, even though t_j has a density, the distribution of j's announcements can have point masses. While virtually the same theory of proper scoring rules applies either to discrete or continuous distributions, to our knowledge there is no theory of proper scoring rules for mixed distributions. Therefore we move to using a scoring rule for the distribution over discrete events.

Lemma 4. *Assumption 4 implies that there exists a $\delta^* > 0$ such that for any $0 < \delta < \delta^*$ and any agent i, there is an agent $j \neq i$ such that T_j contains a finite set of disjoint balls $B^{ij} = \left\{b_1^{ij}, \ldots, b_M^{ij}\right\}$ with radius greater than δ such that for any t_i, t_i' with $t_i \neq t_i'$ there is at least one b_m^{ij} such that $f\left(t_j|t_i\right) \neq f\left(t_j|t_i'\right)$ for all $t_j \in b_m^{ij}$.* [14]

The key distinction between Assumption 4 and Lemma 4 is that Assumption 4 asserts that for any t_i and t_i' there is an open ball (which may depend on t_i and t_i') over which the associated beliefs of different types differ, while Lemma 4 establishes the existence of a finite set of balls such that *no matter which t_i and t_i' are chosen* their associated beliefs differ over at least one ball. To see the role that Lemma 4 will play in the proof of Theorem 2, consider two distinct types t_i and t_i' for agent i. By Lemma 4, let b_m^{ij} be the ball in agent j's announcement space satisfying Lemma 4 that distinguishes these types. If agent j announces truthfully, types t_i and t_i' assign different probabilities to event $t_j \in b_m^{ij}$, and so a scoring rule based on whether agent j's announcement is in b_m^{ij} can be used to truthfully elicit whether agent i's type is t_i or t_i'.

The lower bound on the size of the balls in B^{ij} ensures that there is a partition of events that distinguishes any two types even if j is allowed to distort his announcements using a δ-truthful strategy (for $\delta < \delta^*$). To see how, let b_m^{ij} be the ball in B^{ij} to which t_i and t_i' assign different probabilities to the event $t_j \in b_m^{ij}$. Let r_m be the radius of b_m^{ij}, and let $\hat{b}_m^{ij}$ be the ball with the same center and radius $r_m - \delta$.

If j's strategy is δ-truthful, then the set of types that announce $a_j \in \hat{b}_m^{ij}$ must be contained in b_m^{ij}. Hence whenever $s_j(\cdot) \in C_j(\delta)$, types t_i and t_i', assign different probabilities to the event $a_j \in \hat{b}_m^{ij}$ conditional on $s_j(\cdot)$. To see why, consider Fig. 1, which illustrates the one-dimensional case. Suppose that $b_m^{ij} = \left[t_j^- - \delta, t_j^+ + \delta\right]$. According to Lemma 4, the densities for some t_i's are drawn in such a way that they don't cross over this region. Now, look at the smaller event, $\hat{b}_m^{ij} = \left[t_j^-, t_j^+\right] \subset b_m^{ij}$. Note that since types can only distort their announcements by δ or less, if t_j δ-truthfully announces $a_j \in \hat{b}_m^{ij}$, then $t_j \in b_m^{ij}$. However, since the densities for these values of t_i are ranked over the entire set b_m^{ij}, conditional on $s_j(\cdot) \in C_j(\delta)$, two distinct types whose densities do not cross over b_m^{ij} cannot assign the same probability to j's announcement being in $\hat{b}_m^{ij}$. In Fig. 1, the heavy black lines on the horizontal axis indicate the set of types t_j that make announcements in $\hat{b}_m^{ij}$ for some hypothetical $s_j(\cdot) \in C_j(\delta)$. Looking at the shaded regions above the t_j in this set, the densities for the various values of t_i do not cross. Thus for any two t_i whose densities do not cross over b_{ij}, the one with the higher density must assign higher probability to the event $a_j \in \hat{b}_m^{ij}$ for any possible announcement strategy $s_j(\cdot) \in C_j(\delta)$.

If, as asserted by Uniform Stochastic Relevance and Lemma 4, there is a finite set of balls B^{ij} such that for each possible pair of types there is one ball over which their densities do not cross, then we can use the sets $\hat{b}_1^{ij}, \ldots, \hat{b}_M^{ij}$ along with $\hat{b}_0^{ij} \equiv T_j \backslash \cup \hat{b}_m^{ij}$ (i.e., "everything else") as a partition of events that distinguishes every pair of types for every possible strategy $s_j(\cdot) \in C_j(\delta)$. That is, conditional on $s_j(\cdot)$, different types t_i have different beliefs about the distribution over the events in $\hat{B}^{ij} = \left\{\hat{b}_0^{ij}, \hat{b}_1^{ij}, \ldots, \hat{b}_M^{ij}\right\}$. Thus, transfers based on the quadratic scoring rule applied to

[14] The existence of Lemma 4's B^{ij} clearly implies Uniform Stochastic Relevance. Hence Uniform Stochastic Relevance holds if and only if such a B^{ij} exists.

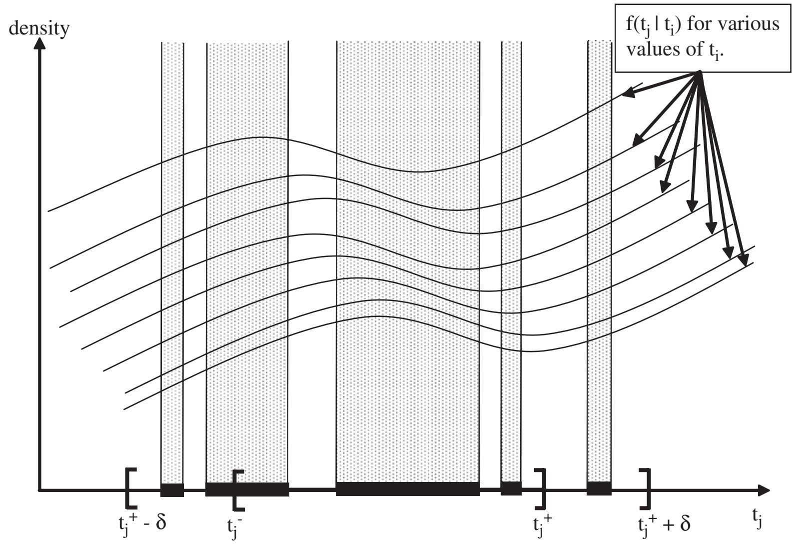

Fig. 1. δ-truthful strategies assign different probabilities to $\left\{a_j \in \left[t_j^-, t_j^+\right]\right\}$.

the events in $\hat{B}^{ij}$ (conditional on $s_j(\cdot)$) are strictly proper. If agent i only cared about the transfer, his best response to $s_j(\cdot)$ under these transfers would be to announce his true type. When agent i also cares about the direct return from the social choice, basing transfers on a sufficiently large scaling of the quadratic scoring rule ensure that agent i's best response is close to truthful.

Theorem 2 establishes that there is a transfer scheme that δ-implements any decision rule in BNE. The main tool employed is Schauder's fixed-point theorem (see [33]). [15]

Theorem 2. *Under Assumptions* 1, 2, *and* 4, *for any decision rule and any* $\delta > 0$ *there exist balanced transfers that* δ*-implement that decision rule in BNE.*

The intuition for the proof begins by noting that each agent's payoff is a linear combination of his direct return and transfer. Thus, the situation where the transfers are multiplied by a large (positive) constant is one where the agent puts small relative weight on his direct return, which is similar to the case where the agent puts zero weight on his direct return. When agents care only about their transfers, transfers based on the quadratic scoring rule ensure that truthful revelation is a strict equilibrium. If we knew that this equilibrium changed smoothly with the relative weight put on agents' direct returns, then games with nearby payoffs would have a nearby equilibrium. Thus, games in which the relative weight on transfers was very high would have an equilibrium in which agents' strategies were nearly truthful. Unfortunately, this smooth dependence property, which is related to lower hemicontinuity of the equilibrium correspondence, does not hold in general. Nevertheless, by exploiting the fact that the truthful equilibrium of the transfers-only game is strict, we are able to show that when agents' strategies are nearly truthful, nearby games satisfy

[15] Meirowitz [27] uses a slightly different version of Schauder's Theorem to prove a general existence result for equilibria in Bayesian games with infinite type and action spaces.

the requirements for the application of Schauder's fixed-point theorem, and thus that games where agents' put small relative weight on their direct returns have a nearly truthful equilibrium. [16]

Theorem 2 establishes the existence of an equilibrium in which agents play nearly truthful strategies. The question remains whether, under the transfers that induce the δ-truthful equilibrium, other equilibria exist as well and, if so, whether those equilibria are also δ-truthful. In general, there is no reason to rule out such equilibria. Given that agents' incentives are primarily driven by their desire to maximize the transfer they receive and that these transfers are determined by how well each agent predicts the announcements of the other agents, it is easy to imagine that there could be equilibria in which all agents permute their announcements in such a way that announcements are no longer close to truthful but still predict other agents' announcements well.

There is, however, one circumstance in which it is possible to establish that all equilibria must be nearly truthful. This is the case in which the center receives a signal of its own that is stochastically related to the agents' types. [17] In effect, for each agent i, the center's information plays the role of the agent j whose information is used to score agent i. Since no agent's behavior can affect the distribution of the center's information, the expected payment from transfers based on the quadratic scoring rule applied to the center's information is uniquely maximized by telling the truth regardless of the other agents' strategies. The argument in Theorem 2 then establishes that for any $\delta > 0$ there exists a δ-truthful equilibrium. Further, since a sufficiently large scaling of the payments ensures that all best responses are δ-truthful, all equilibria must be δ-truthful.

Returning to the case where the center does not receive an informative signal, while Theorem 2 establishes that agents' strategies are nearly truthful, from the perspective of the social planner our real interest is not whether agents are telling the truth, but rather whether the resulting social choice rule is close to that implied by the planner's desired rule, and whether realized social welfare is close to the planner's desired welfare level. These desirable properties follow, however, because transfers are balanced and agents' payoffs are assumed to be continuous conditional on the social choice function (Assumption 1).

MR shows that when agents' types are correlated, for any game the center can extract from each agent nearly all of the rents that agent earns by participating in the game. [18] Their mechanism offers agents a finite menu of participation fee schedules such that, when the agent selects his preferred schedule and then plays the game, he is left with a rent that, though positive, is arbitrarily small.

While MR shows that *for a given game,* a participation fee schedule can ensure that agents' interim participation constraints can be satisfied at (nearly) no cost to the center, they do not address the question of whether, for a given decision rule, a game exists that implements that decision rule. In particular, the center cannot extract the full information rent (i.e., the rent that would be generated if the center could observe agents' types and make ex post efficient decision) unless there exists a mechanism that implements the ex post efficient decision rule. Prior to this paper, there have been no results that show the general existence in the standard mechanism design framework of an ex post efficient mechanism when agents have multidimensional, continuous types and interdependent valuations. [19] To the extent that ex post efficient mechanisms have been

[16] The key step is to establish that agents' best responses to any δ-truthful opponents' strategies are unique, which requires conditioning transfers on the other agents' strategies.

[17] If the center's information has a density, then the center's information must satisfy the player j role in Assumption 3.

[18] Their required condition (*) is strictly stronger than our Assumption 3.

[19] As discussed earlier, Mezzetti's [28] mechanism operates in a slightly different framework than the standard models and uses a weaker form of budget balance.

shown to exist, these results typically require additional assumptions on the form of agents' direct returns functions, e.g., single crossing. The results in this paper do not impose any restrictions on direct returns functions beyond smoothness (Assumption 1).

We show that if agents' types are correlated, then any decision rule, including the ex post efficient decision rule, can be implemented arbitrarily closely (i.e., δ-truthfully). Provided that beliefs satisfy MR's condition (*), our result, coupled with the MR result, establishes that the center can extract (approximately) the full information rent and satisfy agents' interim participation constraints by first offering agents a menu of participation fees and then running our scoring-rule based system.

5. Limitations of quasilinear mechanism design

This paper employs the quasilinear mechanism design framework, and as such it suffers from the well-known limitations of the approach. [20] These include, first, that the transfers needed to induce (near) truth-telling may be very large, and thus for small δ our mechanism may be infeasible if agents face limited liability constraints. Second, the quasilinear framework assumes that agents are risk neutral with respect to their transfers. If agents are risk averse over their transfers, then it will not generally be possible to (nearly) implement any decision rule with budget balance. However, if the center is interested in inducing (nearly) truthful revelation without budget balance, then redefining the transfers in terms of utilities instead of monetary amounts will accomplish this goal.

Recently, Neeman [29] and Heifetz and Neeman [15] have launched another line of criticism against the literature on mechanism design with correlated information. They argue that although the correlation requirements employed in the literature appear rather reasonable, they have the common feature that an agent's beliefs uniquely determine his preferences, which they term the BDP property. Stochastic relevance, as embodied in Assumptions 3 and 4 of this paper, implies the BDP property. Heifetz and Neeman [15] show that the BDP property is a non-generic property of the universal type space. Thus, while correlation seems like a reasonable assumption, the set of BDP beliefs is, in a sense, "small."

Another line of criticism regarding Bayesian mechanism design is that Bayesian equilibrium is belief-based. As such, incentive compatible mechanisms are highly sensitive to the information structure of the problem. MR observes that such dependence casts doubt on whether such results teach us much about real-world asymmetric information problems. This criticism has led to the search for "robust" mechanisms that do not depend on agents' beliefs about others' information, usually employing the concept of ex post equilibrium. In light of this, the JM result on the generic impossibility of efficient BIC design with independently distributed types also implies the impossibility of ex post incentive compatibility. Our results provide a counterpoint to JM by showing that BIC design is possible in their environment if the independence assumption is relaxed. However, our BIC mechanism is not ex post incentive compatible. Indeed, Jehiel et al. [17] show that only constant decision rules are implementable in ex post equilibrium in generic mechanism design problems with multidimensional, continuous types and interdependent valuations.

---

[20] Crémer and McLean [6] discuss the limited liability and risk neutrality assumptions in the context of their full extraction result.

6. Conclusion

This paper extends the mechanism design literature to show that when agents' types are continuous, multidimensional, and mutually payoff relevant, incentive-compatible implementation of any decision rule is possible provided that agents' types satisfy one of our rather mild correlation conditions. Thus we provide a complement to the JM impossibility result for the case of independent information. Our results also complement MR by showing that there is an ex post efficient mechanism in the multidimensional, continuous, mutually payoff-relevant case.

While we show the existence of transfers that induce a δ-truthful BNE, we have not considered the question of whether there exist transfers that render the exact truth a BNE. This is a technically daunting task; it remains an open question.

The scoring-rule based approach we adopt has the advantage of being simpler than the approaches commonly adopted in the mechanism design literature. Stochastic relevance (as embodied in Assumption 3 or 4) requires verifying only that distributions are different for different types, which is substantially easier than verifying the compatibility condition of d'Aspremont and Gérard-Varet [7,8], the linear independence condition of Crémer and McLean [5,6], or the generalization of the Crémer–McLean condition found in MR, each of which must hold for all prior distributions for each agent's type. Beyond its advantage of simplicity, stochastic relevance is also slightly weaker than any of these conditions. [21] The scoring-rule-based payments used in our mechanism are also relatively easy to construct and our proofs provide a blueprint for doing so. Our approach represents an advance over existing methods, which generally prove the existence of a mechanism but provide little or no guidance on how it can be constructed. [22]

Acknowledgments

We thank Drew Fudenberg, Jerry Green, Daniel Hojman, Deb Hughes-Hallett, Matt Jackson, John H. Lindsey II, Rich McLean, Zvika Neeman, Andy Postlewaite, Phil Reny, Bill Sandholm, Michael Schwarz, Ed Shpiz, various seminar participants, an associate editor and two anonymous referees for helpful comments. Zeckhauser gratefully acknowledges financial support of NSF Grant IIS-0428868.

Appendix A. Proofs

Proof of Lemma 1. Suppose not. Then there exists $\delta > 0$ such that for all n there exists a sequence of pairs of types $\left(t_{in}, t'_{in}\right)$ such that $\left|\left|t_{in} - t'_{in}\right|\right|_R \geqslant \delta$ and $\left\| f_j\left(\cdot|t_{in}\right) - f_j\left(\cdot|t'_{in}\right)\right\|_2 \leqslant 1/n$. By compactness, this sequence has a convergent subsequence. Let the limit be $\left(t_{i*}, t'_{i*}\right)$, and note that $t_{i*} \neq t'_{i*}$. We have that $\lim \left\| f_j\left(\cdot|t_{i*}\right) - f_j\left(\cdot|t'_{i*}\right)\right\|_2 = 0$, and

$$\lim \left\| f_j\left(\cdot|t_{in}\right) - f_j\left(\cdot|t'_{in}\right)\right\|_2 = \lim \left(\int \left|f_j\left(t_j|t_{in}\right) - f_j\left(t_j|t'_{in}\right)\right|^2 dt_j\right)^{1/2}$$

[21] To be fair, the task of full surplus extraction is more demanding than (nearly) truthful implementation, and so while these papers employ stricter conditions they also achieve stronger results. Our condition is very similar to that employed by Aoyagi [1] in the finite case.

[22] Frequently, such approaches rely on a linear systems approach to demonstrate existence. See [9] for a survey of the use of this method.

$$\geqslant \lim \left(\int \min \left\{ 1, \left| f_j \left(t_j | t_{in} \right) - f_j \left(t_j | t'_{in} \right) \right|^2 \right\} dt_j \right)^{1/2}$$
$$= \left(\int \lim \left(\min \left\{ 1, \left| f_j \left(t_j | t_{in} \right) - f_j \left(t_j | t'_{in} \right) \right|^2 \right\} \right) dt_j \right)^{1/2}$$
$$= \left(\int \min \left\{ 1, \lim \left| f_j \left(t_j | t_{in} \right) - f_j \left(t_j | t'_{in} \right) \right|^2 \right\} dt_j \right)^{1/2}$$
$$= \left(\int \min \left\{ 1, \left| f_j \left(t_j | t_{i*} \right) - f_j \left(t_j | t'_{i*} \right) \right|^2 \right\} dt_j \right)^{1/2},$$

where the fact that $\min \left\{ 1, \left| f_j \left(t_j | t_{in} \right) - f_j \left(t_j | t'_{in} \right) \right|^2 \right\} \leqslant 1$ allows us to apply the Bounded Convergence Theorem [30, p. 84] in moving from the second to third line. Since

$$\left(\int \min \left\{ 1, \left| f_j \left(t_j | t_{i*} \right) - f_j \left(t_j | t'_{i*} \right) \right|^2 \right\} dt_j \right)^{1/2} = 0,$$

$\left| f \left(t_j | t_{i*} \right) - f \left(t_j | t'_{i*} \right) \right| = 0$ for almost all t_j. Since $f \left(t_j | t_i \right)$ is continuous, $f \left(t_j | t_{i*} \right) = f \left(t_j | t'_{i*} \right)$ for all t_j. However, this violates Assumption 3. □

Proof of Lemma 2. The result is standard. This proof follows [31]. Let $\Upsilon (a_i | t_i) = \int Q \left(t_j | a_i \right) f \left(t_j | t_i \right) dt_j$ be agent i's expected transfer when $s_j (\cdot) = \tau_j$. Substituting in the definition of the quadratic scoring rule, we have

$$\Upsilon (a_i | t_i) = \int_{T_j} \left(2 f \left(t_j | a_i \right) - \int_{T_j} f \left(t_j | a_i \right)^2 dt_j \right) f \left(t_j | t_i \right) dt_j .$$

Rearranging $\Upsilon (a_i | t_i)$ yields:

$$\Upsilon (a_i | t_i) = \int_{T_j} f \left(t_j | t_i \right)^2 dt_j - \int_{T_j} \left(f \left(t_j | a_i \right) - f \left(t_j | t_i \right) \right)^2 dt_j .$$

Hence:

$$\Upsilon (a_i | t_i) - \Upsilon (t_i | t_i) = - \int_{T_j} \left(f \left(t_j | a_i \right) - f \left(t_j | t_i \right) \right)^2 dt_j ,$$

which is zero when $a_i = t_i$ and strictly negative otherwise. □

Proof of Lemma 3. From Lemma 1, for any δ there exists a $\mu > 0$ such that $\left\| f \left(t_j | a_i \right) - f \left(t_j | t_i \right) \right\|_2 \geqslant \mu$ for all a_i and t_i with $\| a_i - t_i \| \geqslant \delta$. Using the notation from the proof of Lemma 2,

$$\begin{aligned} &\left| \int_{T_j} Q \left(t_j | a_i \right) f \left(t_j | t_i \right) dt_j - \int_{T_j} Q \left(t_j | a_i \right) f \left(t_j | a_i \right) dt_j \right| \\ &\quad = \left| \Upsilon (a_i | t_i) - \Upsilon (t_i | t_i) \right| \\ &\quad = \left| - \int_{T_j} \left(f \left(t_j | a_i \right) - f \left(t_j | t_i \right) \right)^2 dt_j \right| \\ &\quad = \left(\left\| f \left(t_j | a_i \right) - f \left(t_j | t_i \right) \right\|_2 \right)^2 \geqslant \mu^2 . \end{aligned}$$

Letting $\varepsilon = \mu^2$ completes the proof. □

Proof of Theorem 1. Consider agent i, and suppose all other agents announce truthfully, $s_j(\cdot) = \tau_j, \forall j \neq i$. Thus for the remainder of the proof we can replace a_j with t_j in the expression for agent i's payoff. Since expected direct returns are continuous in announcements and $\times T_i$ is compact, expected direct returns are uniformly continuous. Hence for any $\varepsilon > 0$ there exists $\delta > 0$ such that $||t_i - a_i||_R \leqslant \delta$ implies $|E(v_i(t, a_i, t_{-i})|t_i) - E(v_i(t, t_i, t_{-i})|t_i)| \leqslant \varepsilon$.

Given δ, let $T_i' = \{(t_i, a_i) \in T_i \times T_i : ||t_i - a_i||_R \geqslant \delta\}$ be the set of type-announcement pairs that are at least δ apart. For each i, choose $j(i)$ according to Assumption 2. Define the payments to agent i according to the quadratic scoring rule:

$$x_i(a_i, t_{-i}) = 2f\left(t_{j(i)}|a_i\right) - \int_{T_{j(i)}} f\left(t_{j(i)}|a_i\right)^2 dt_{j(i)}. \tag{1}$$

Since the expected quadratic score is uniquely maximized at $a_i = t_i$ and T_i' is compact, by Lemma 3 there exists an $\hat{\varepsilon} > 0$ such that for any $(t_i, a_i) \in T_i'$:

$$E\left\{x_i\left(a_i, t_{j(i)}\right)|t_i\right\} \leqslant E\left\{x_i\left(t_i, t_{j(i)}\right)|t_i\right\} - \hat{\varepsilon}. \tag{2}$$

Next, scale the payments to agent i according to $x_i^*\left(a_i, t_{j(i)}\right)$:

$$x_i^*\left(a_i, t_{j(i)}\right) = \frac{\left(2\bar{M} + 1\right)}{\hat{\varepsilon}} x_i\left(a_i, t_{j(i)}\right). \tag{3}$$

Consider the expected utility reaped by a truthful announcement as compared to announcing a_i with $(t_i, a_i) \in T_i'$:

$$\begin{aligned}
&E\left\{v_i(t, a_i, t_{-i}) + x_i^*\left(a_i, t_{j(i)}\right)|t_i\right\} - E\left\{v_i(t, t_i, t_{-i}) + x_i^*\left(t_i, t_{j(i)}\right)|t_i\right\} \\
&\quad = E\left\{v_i(t, a_i, t_{-i}) - v_i(t, t_i, t_{-i})|t_i\right\} + E\left\{x_i^*\left(a_i, t_{j(i)}\right) - x_i^*\left(t_i, t_{j(i)}\right)|t_i\right\} \\
&\quad < 2\bar{M} + \left(2\bar{M} + 1\right)\left[\frac{\left[x_i\left(a_i, t_{j(i)}\right)\right]}{\hat{\varepsilon}} - \frac{\left[x_i\left(t_i, t_{j(i)}\right)\right]}{\hat{\varepsilon}}\right] \\
&\quad < 2\bar{M} + \left(2\bar{M} + 1\right)\left[-\frac{\hat{\varepsilon}}{\hat{\varepsilon}}\right] < 0.
\end{aligned}$$

Hence under payment scheme $x_i^*\left(a_i, t_{j(i)}\right)$, announcing truthfully earns a higher payoff than any announcement such that $(a_i, t_i) \in T_i'$. And, by the choice of δ, announcements a_i such that $(a_i, t_i) \notin T_i'$ have lower expected transfers than truthful announcements and have expected direct returns that exceed those of truthful announcement by less than ε. Hence truthful announcement is an ε-best response. Since i is chosen arbitrarily, payments can be constructed that make the truth an ε-best response for all agents.

To balance the budget, for each agent i choose an agent $\varkappa(i) \notin \{i, j(i)\}$ with the understanding that $\varkappa(i)$ will fund i's transfer. Let $K_i = \{k|\varkappa(k) = i\}$ be the set of all agents whose transfers i funds. Agent i's net transfer is therefore:

$$x_i(a_i, t_{-i}) = x_i^*\left(a_i, t_{j(i)}\right) - \sum_{k \in K_i} x_k^*\left(a_k, t_{j(k)}\right). \tag{4}$$

Since agent i's announcement does not affect the terms after the summation, his incentives are not affected, and transfer scheme $x(t, a)$ is ε-BIC and balances the budget. □

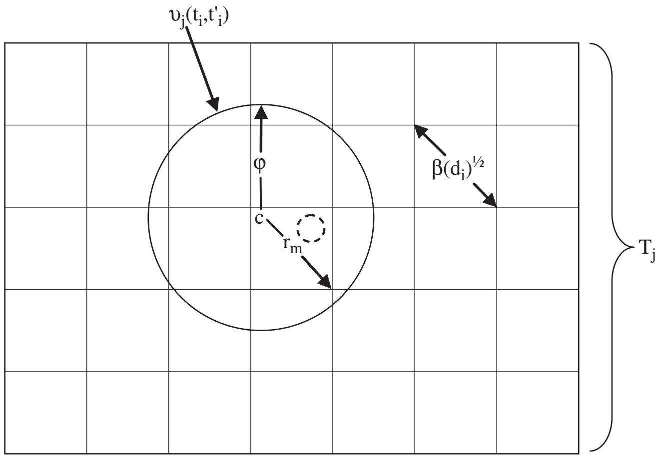

Fig. A.1. Assumption 4 implies Lemma 4.

Proof of Lemma 4. For each i choose an appropriate j according to Assumption 4. Lay a d_i-dimensional rectangular grid over T_j by dividing each of the d_i dimensions into increments of size $\beta > 0$. Thus, the grid divides T_j into hypercubes with sides of length β. The maximum distance between any two points in a hypercube is $\beta\sqrt{d_i}$ (i.e., the distance in $\mathbb{R}^{d_i}$ between $(0, 0, \ldots, 0)$ and $(\beta, \beta, \ldots, \beta)$). Choose β such that $\beta\sqrt{d_i} < \phi$. Consider two distinct types t_i and t_i' for agent i. By Assumption 4, for any t_i and t_i', there exits a ball $\theta_j\left(t_i, t_i'\right)$ with radius ϕ such that $f\left(t_j|t_i\right) \neq f\left(t_j|t_i'\right)$ for all $t_j \in \theta_j\left(t_i, t_i'\right)$. Let c be the center of $\theta_j\left(t_i, t_i'\right)$. As illustrated in Fig. A.1, the maximum distance between c and any point in the same hypercube is $r_m \leqslant \beta\sqrt{d_i} < \phi$, and thus the hypercube containing c is contained in $\theta_j\left(t_i, t_i'\right)$.[23] Since the side length of each hypercube is β, there is a ball of radius $\beta/3$ within this hypercube (illustrated by the dotted circle) such that $f\left(t_j, t_i\right) \neq f\left(t_j, t_i'\right)$ for all t_j within the ball. Since there are a finite set of grid elements, taking one such ball for each grid element defines a finite set of disjoint balls B^{ij} such that for each distinct t_i and t_i' there is at least one ball on which $f\left(t_j|t_i\right)$ and $f\left(t_j|t_i'\right)$ are not equal. □

Proof of Theorem 2. The proof employs transfers based on a large scaling of the quadratic scoring rule. However, rather than working with h_i as the (large) scaling applied to the transfers, yielding payoffs $u_i(a, t) = v_i(t, a) + h_i x_i\left(a_i, a_j\right)$, we will instead work with the equivalent formulation in which $\gamma_i = 1/h_i$ and payoffs are given by $\tilde{u}_i(a, t) = \gamma_i v_i(t, a) + x_i\left(a_i, a_j\right)$. In this formulation, the agent's utility function depends continuously on γ_i. The game with $\gamma_i = 0$ is one in which agent i cares only about the transfer, and the game with γ_i positive but small (which corresponds a large scaling of the transfers) can be thought of as a slightly perturbed version of the $\gamma_i = 0$ game. The proof exploits the fact that if the conditions for application of Schauder's fixed-point theorem are satisfied when $\gamma_i = 0$ and best response correspondences are single-valued and suitably continuous, then they are also satisfied for small-but-positive values of γ_i.

[23] For c on the boundary between several hypercubes, all such hypercubes are contained in $\theta_j\left(t_i, t_i'\right)$, and any such hypercube can be used for the remainder of the argument.

Without loss of generality, assume $\delta \leqslant \delta^*$ as specified in Lemma 4.[24] For each i choose a player j satisfying Lemma 4, and let $\hat{B}^{ij} = \left\{ \hat{b}^{ij}_0, \hat{b}^{ij}_1, \ldots, \hat{b}^{ij}_M \right\}$, denote the partition of agent j's announcement space described above.[25] Let

$$p^{s_j(\cdot)}\left(\hat{b}^{ij}|a_i\right) \equiv \int_{\left\{t_j|s_j(t_j)\in\hat{b}^{ij}\right\}} f\left(t_j|a_i\right) dt_j,$$

be the probability of event $\hat{b}^{ij}$ if agent i's type is a_i, conditional on j's announcement strategy. If agent j plays strategy $s_j(\cdot)$, let the transfer to agent i be $x_i^{s_j(\cdot)}\left(a_i, a_j\right)$, which is based on the quadratic scoring rule applied to the events in $\hat{B}^{ij}$ according to:

$$x_i^{s_j(\cdot)}\left(a_i, a_j\right) = 2p^{s_j(\cdot)}\left(\hat{b}^{ij}\left(a_j\right)|a_i\right) - \sum_{m=1}^{M} p^{s_j(\cdot)}\left(\hat{b}^{ij}_m|a_i\right)^2.$$

Note that for any $s_j(\cdot) \in C_j(\delta)$, by Lemma 4, transfers $x_i^{s_j(\cdot)}\left(a_i, a_j\right)$ represent a strictly proper scoring rule, and hence for each, $i, s_j(\cdot) \in C_j(\delta)$, and t_i, agent i's expected transfer is maximized by announcing truthfully ($a_i = t_i$).

For any $s_j(\cdot)$, for each value of t_i agent i chooses a_i to maximize:

$$\gamma_i \int_{T_{-i}} v_i\left(t, a_i, s_{-i}(t_{-i})\right) f\left(t_{-i}|t_i\right) dt_{-i} + \int_{T_j} x_i^{s_j(\cdot)}\left(a_i, a_j\right) f\left(t_j|t_i\right) dt_j. \tag{5}$$

Since $p^{s_j(\cdot)}\left(\hat{b}^{ij}|a_i\right)$ is continuous in $s_j(\cdot)$ (which implies that $x_i^{s_j(\cdot)}\left(a_i, a_j\right)$ is continuous in $s_j(\cdot)$) and $v_i(t, a)$ is continuous in t and a, (5) is continuous in $s_{-i}(\cdot)$, and in particular in $s_j(\cdot)$. Further, our assumptions ensure that (5) is continuously differentiable in a_i. Note that the transfers are constructed for each $s_j(\cdot)$ in order to render truthful reporting a best response when $\gamma_i = 0$ (and thus, as we show below, they render nearly truthful reporting a best response when γ_i is positive but sufficiently small).

Let $BA_i\left(t_i, s_{-i}(\cdot), \gamma_i\right)$ denote agent i's best announcements, i.e., announcements that maximize (5) given his type, the other agents' strategies, and the weight i places on his direct return. $BA_i\left(t_i, s_{-i}(\cdot), \gamma_i\right)$ may be multi-valued. Let $BR_i\left(s_{-i}(\cdot), \gamma_i\right)$ denote the best response operator for agent i, parameterized by γ_i. That is, for a given value of γ_i, $BR_i\left(s_{-i}(\cdot), \gamma_i\right)$ maps $s_{-i}(\cdot)$ to best-action correspondences $BA_i\left(t_i, s_{-i}(\cdot), \gamma_i\right)$. Let $BR\left(s(\cdot), \gamma\right) = \left(BR_1\left(s_{-1}(\cdot), \gamma_1\right), \ldots, BR_I\left(s_{-I}(\cdot), \gamma_I\right)\right)$ denote the best response operator for all agents, where $\gamma = \left(\gamma_1, \ldots, \gamma_I\right)$. Let $C_i^{\mathrm{E}}(\delta)$ be the uniformly equicontinuous set of δ-truthful strategies for player i defined by

$$\begin{aligned} C_i^{\mathrm{E}}(\delta) = \Big\{ & s_i(\cdot) \in C_i(\delta) : \forall \psi > 0, \forall t_i, t_i' \in T_i, \left\| t_i' - t_i'' \right\|_R < \psi \\ & \text{implies } \left\| s_i\left(t_i'\right) - s_i\left(t_i''\right) \right\|_R < 2\psi \Big\}. \end{aligned}$$

We define $C_{-i}^{\mathrm{E}}(\delta)$ and $C^{\mathrm{E}}(\delta)$ in the usual way. Since $C_i^{\mathrm{E}}(\delta)$ is uniformly bounded, closed, and equicontinuous, the Arzela–Ascoli Theorem implies that it is compact [21, p. 273]. The remainder of the proof shows that for γ sufficiently small, $BR\left(s(\cdot), \gamma\right)$ is a continuous map from $C_i^{\mathrm{E}}(\delta)$

[24] If not, use δ^* in the following construction. Since it establishes that there is a δ^*-truthful BNE, this BNE is also δ-truthful for $\delta > \delta^*$.

[25] We denote the player as j rather than $j(i)$ for notational convenience.

into itself, and thus by Schauder's fixed-point theorem has a fixed point. Such a fixed point is a δ-truthful BNE of the announcement game. [26]

Step 1: Note that for each i, T_i is a compact, convex subset of a finite-dimensional Euclidean space with a non-empty interior, and that agent i's objective function (5) is jointly continuous in $t_i, a_i, s_{-i}(\cdot)$, and γ_i. Thus for any t_i, γ_i, and $s_{-i}(\cdot) \in C^{\mathrm{E}}_{-i}(\delta)$, agent i's best response exists and by the Theorem of the Maximum [2] $BA_i\left(t_i, s_{-i}(\cdot), \gamma_i\right)$ is upper hemicontinuous in t_i and γ_i.

Step 2: Next, we show that for γ_i sufficiently small, for any $s_{-i}(\cdot) \in C^{\mathrm{E}}_{-i}(\delta)$ and any $a_i \in BA_i\left(t_i, s_{-i}(\cdot), \gamma_i\right)$, $\|a_i - t_i\|_R \leqslant \delta$. Suppose that for some $s_{-i}(\cdot)$ and t_i, for any $\gamma_i > 0$ there exists $a_i \in BA_i\left(t_i, s_{-i}(\cdot), \gamma_i\right)$ and $\|a_i - t_i\| > \delta$. Take a sequence $\gamma_i = 1/n$, and let $a_i^n \in BA_i\left(t_i, s_{-i}(\cdot), 1/n\right)$. By compactness, a_i^n has a convergent subsequence. Let a_i^* be the limit point, and note that $\left\|a_i^* - t_i\right\| \geqslant \delta$. Let $W_i^{s_j(\cdot)}(a_i, t_i) = \int_{T_j} x_i^{s_j(\cdot)}\left(a_i, a_j\right) f\left(t_j|t_i\right) dt_j$. Since $BA_i\left(t_i, s_{-i}(\cdot), \gamma_i\right)$ is upper hemicontinuous, $a_i^* \in BA_i\left(t_i, s_{-i}(\cdot), 0\right)$, which contradicts that $a_i = t_i$ is the unique maximizer of $W_i^{s_j(\cdot)}(a_i, t_i)$. Hence for γ_i sufficiently small, all best responses are within δ of being truthful.

Step 3: Next, we argue that for γ_i sufficiently small, $BA_i\left(t_i, s_{-i}(\cdot), \gamma_i\right)$ is single-valued. Fix $s_{-i}(\cdot)$. Since these transfers are strictly proper, for any t_i, $a_i = t_i$ maximizes $W_i^{s_j(\cdot)}(a_i, t_i)$ strictly and uniquely. Since $a_i = t_i$ maximizes $W_i^{s_j(\cdot)}(a_i, t_i)$, $a_i = t_i$ satisfies the first-order necessary conditions: $D_{a_i} W_i^{s_j(\cdot)}(t_i, t_i) = \bar{0}$, where D_{a_i} denotes the gradient vector with respect to the components of a_i and $\bar{0}$ denotes the d_i-dimensional zero (column) vector. [27] This implies that $W_i^{s_j(\cdot)}(a_i, t_i) < W_i^{s_j(\cdot)}(t_i, t_i) + D_{a_i} W_i^{s_j(\cdot)}(t_i, t_i) \cdot (a_i - t_i)$, which establishes that $W_i^{s_j(\cdot)}(a_i, t_i)$ is locally strictly concave in a_i at $a_i = t_i$. By continuity, $W_i^{s_j(\cdot)}(a_i, t_i)$ is locally strictly concave in a_i for a_i near t_i, as is $W_i^{s_j(\cdot)}(a_i, t_i) + \gamma_i E\left(v_i(a, t)\,|s_j(\cdot), t_i\right)$ for γ_i sufficiently small (since $v_i(a, t)$ has bounded first and second derivatives). Let $\rho_i^{s_j(\cdot)} > 0$ be the largest ρ such that $W_i^{s_j(\cdot)}(a_i, t_i)$ is locally strictly concave in a_i for all t_i and a_i in the intersection of T_i and the open ball with center at t_i and radius ρ. Let ρ_i^* be the minimum of all $\rho_i^{s_j(\cdot)}$. To establish that ρ_i^* is strictly positive, suppose not. In this case there exists a sequence of strategies $s_j^n(\cdot) \in C_j^{\mathrm{E}}(\delta)$ such that $\rho_i^{s_j^n(\cdot)} \to 0$. Since $C_j^{\mathrm{E}}(\delta)$ is compact, this sequence has a convergent subsequence. Let $s_j^*(\cdot) \in C_j^{\mathrm{E}}(\delta)$ denote the limit point. We have then that there exists t_i such that $W_i^{s_j^*(\cdot)}(a_i, t_i) \geqslant W_i^{s_j^*(\cdot)}(t_i, t_i)$ for all a_i in a neighborhood of t_i. However, this contradicts that $a_i = t_i$ is the unique maximizer of $W_i^{s_j^*(\cdot)}(a_i, t_i)$ at t_i. Hence $\rho_i^* > 0$.

By continuity, for γ_i sufficiently small, $W_i^{s_j(\cdot)}(a_i, t_i) + \gamma_i E\left(v_i(a, t)\,|s_j(\cdot), t_i\right)$ is locally strictly concave for all t_i and all a_i with $\|a_i - t_i\| \leqslant \rho_i^*/2$. The argument of the previous paragraph establishes that for γ_i sufficiently small, all best responses are within $\rho_i^*/2$, and hence for γ_i sufficiently small that $BA_i\left(t_i, s_{-i}(\cdot), \gamma_i\right)$ is single valued.

---

[26] Using a slightly different version of Schauder's theorem, Meirowitz [27] proves the existence of a BNE in general Bayesian games with infinite type and actions spaces.

[27] This also holds true for t_i on the boundary of T_i since $a_i = t_i$ is the unique global maximizer of $W_i^{s_j(\cdot)}(a_i, t_i)$, and would be even if t_i were not constrained to come from T_i. Since the relevant functions are defined over an open set containing T_i^2 (see footnote 9), it must be that $D_{a_i} W_i^{s_j(\cdot)}(t_i, t_i) = \bar{0}$ for t_i on the boundary of T_i.

Step 4: Next, we show that there exists $\gamma_i^* > 0$ such that $BR_i\left(s_{-i}(\cdot), \gamma_i\right)$ is a continuous map from $C_{-i}^{\mathrm{E}}(\delta)$ to $C_i(\delta)$. Let γ_i^* be the largest γ_i such that $BA_i\left(t_i, s_{-i}(\cdot), \gamma_i\right)$ is single valued for all $s_{-i}(\cdot) \in C_{-i}^{\mathrm{E}}(\delta)$ and all t_i. Compactness of $C_{-i}^{\mathrm{E}}(\delta)$ ensures that $\gamma_i^* > 0$. If not, then there exists a sequence $\left(\gamma_n, t_i^n, s_{-i}^n(\cdot)\right)$ with $\gamma_n \to 0$ such that $BA_i\left(t_i^n, s_{-i}^{\gamma_n}(\cdot), \gamma_n\right)$ is multi-valued. Hence $\left\|BA_i\left(t_i^n, s_{-i}^{\gamma_n}(\cdot), \gamma_n\right) - t_i\right\| > \rho_i^*/2$. By compactness, we can without loss of generality assume that $\left(\gamma_n, t_i^n, s_{-i}^n(\cdot)\right)$ converges. Denote the limit $\left(t_i^*, s_{-i}^*(\cdot), 0\right)$. Since $BA_i\left(t_i, s_{-i}(\cdot), \gamma_i\right)$ is upper hemicontinuous, we have that $\lim BA_i\left(t_i^n, s_{-i}^n(\cdot), \gamma_n\right) = BA_i\left(t_i^*, s_{-i}^*(\cdot), 0\right)$, and thus that for any $a_i \in BA_i\left(t_i^*, s_{-i}^*(\cdot), 0\right)$, $\left\|a_i - t_i^*\right\| > \rho_i^*/2$. However, this contradicts the conclusion of the previous paragraph.

For the remainder of the proof, assume that $\gamma_i < \gamma_i^*$. Since $\gamma_i < \gamma^*$, $BA_i\left(t_i, s_{-i}(\cdot), \gamma_i\right)$ is upper hemicontinuous and single-valued in t_i for any $s_{-i}(\cdot) \in C_{-i}^{\mathrm{E}}(\delta)$, this implies that $BA_i\left(t_i, s_{-i}(\cdot), \gamma_i\right)$ is a continuous function of t_i, and hence that $BR_i\left(s_{-i}(\cdot), \gamma_i\right)$ maps to continuous functions of t_i for γ_i sufficiently small: $BR_i\left(s_{-i}(\cdot), \gamma_i\right) : C_{-i}^{\mathrm{E}}(\delta) \to C_i(\delta)$. Further, since (5) depends continuously on $s_{-i}(\cdot)$ and γ_i, the Theorem of the Maximum also establishes that $BR_i\left(s_{-i}(\cdot), \gamma_i\right)$ is continuous in $s_{-i}(\cdot)$ and γ_i (since $BR_i\left(s_{-i}(\cdot), \gamma_i\right)$ is upper hemicontinuous and single-valued).

Step 5: Finally, we argue that for γ_i sufficiently small, $BR_i\left(s_{-i}(\cdot), \gamma_i\right)$ maps $C_{-i}^{\mathrm{E}}(\delta)$ to $C_i^{\mathrm{E}}(\delta)$. We use $BR_i\left(C_{-i}^{\mathrm{E}}(\delta), \gamma_i\right)$ to denote the set of strategies that are best responses to some strategy in $C_{-i}^{\mathrm{E}}(\delta)$. To establish equicontinuity, we must show that for any $\psi > 0$, and any $t_i, t_i' \in T_i$:

$$\left\|t_i' - t_i''\right\|_R < \psi \quad \text{implies} \sup_{s_i(\cdot)\in BR_i\left(C_{-i}^{\mathrm{E}}(\delta),\gamma_i\right)} \left\|s_i\left(t_i'\right) - s_i\left(t_i''\right)\right\|_R < 2\psi. \tag{6}$$

However, note that $BA_i(t_i, s_{-i}(\cdot), 0) = t_i$ for any $s_{-i}(\cdot) \in C_{-i}^{\mathrm{E}}(\delta)$, and thus that:

$$\left\|t_i' - t_i''\right\|_R < \psi \quad \text{implies} \sup_{s_i(\cdot)\in BR_i\left(C_{-i}^{\mathrm{E}}(\delta),0\right)} \left\|s_i\left(t_i'\right) - s_i\left(t_i''\right)\right\|_R < \psi. \tag{7}$$

Hence (6) is satisfied when $\gamma_i = 0$. Since $BR_i\left(s_{-i}(\cdot), \gamma_i\right)$ is continuous in γ_i, (6) is also satisfied for γ_i sufficiently small. Let γ_i^{**} be such that (6) holds for $\gamma_i < \gamma_i^{**}$. A compactness argument similar to the one used above to show that $\gamma_i^* > 0$ establishes that $\gamma_i^{**} > 0$.

For the remainder of the proof, consider only $\gamma_i \leqslant \gamma_i^{***} = \min\left\{\gamma_i^*, \gamma_i^{**}\right\}$. For such γ_i, $BR_i\left(s_{-i}(\cdot), \gamma_i\right)$ maps $C_{-i}^{\mathrm{E}}(\delta)$ to $C_i^{\mathrm{E}}(\delta)$.

Step 6: Next, we apply Schauder's fixed-point theorem. Schauder's fixed-point theorem says that a continuous operator that maps a nonempty, compact, convex subset of a Banach space into itself has a fixed point [33, Corollary 2.13]. The preceding argument establishes that $C^{\mathrm{E}}(\delta)$ is such a subset and that $BR(s(\cdot), \gamma)$ is a continuous operator that maps $C^{\mathrm{E}}(\delta)$ into itself when $\gamma_i \leqslant \gamma_i^{***}$ for all i. Hence, $BR(s(\cdot), \gamma)$ has a fixed point. The fixed point of the best-response mapping is a δ-truthful BNE of this game, and thus a δ-truthful BNE of the game where payoffs are given by $u_i(a, t) = v_i(t, a) + h_i x_i\left(a_i, a_j\right)$ for $h_i = 1/\gamma_i$. Transfers are balanced using the same type of permutation as was employed in the proof of Theorem 1. □

References

[1] M. Aoyagi, Correlated types and Bayesian incentive compatible mechanisms with budget balancedness, J. Econ. Theory 79 (1998) 142–151.

[2] C. Berge, Topological Spaces: Including a Treatment of Multi-valued Functions, Vector Spaces, and Convexity, Dover Publications, Inc., Mineola, NY, 1998, Translated by E.M. Patterson.
[3] E. Clarke, Multipart pricing of public goods, Public Choice 8 (1971) 19–33.
[4] R.M. Cooke, Experts in Uncertainty: Opinion and Subjective Probability in Science, Oxford University Press, New York, 1991.
[5] J. Crémer, Optimal selling strategies under uncertainty for a discriminating monopolist when demands are interdependent, Econometrica 53 (1985) 345–361.
[6] J. Crémer, Full extraction of surplus in Bayesian and dominant strategy auctions, Econometrica 56 (1988) 1247–1257.
[7] C. d'Aspremont, L.-A. Gérard-Varet, Incentives and incomplete information, J. Public Econ. 11 (1979) 25–45.
[8] C. d'Aspremont, L.-A. Gérard-Varet, Bayesian incentive compatible beliefs, J. Math. Econ. 10 (1982) 83–103.
[9] C. d'Aspremont, L.-A. Gérard-Varet, Linear inequality methods to enforce partnerships under uncertainty: an overview, Games Econ. Behav. 25 (1998) 311–336.
[10] C. d'Aspremont, J. Crémer, L.-A. Gérard-Varet, Incentives and the existence of Pareto-optimal revelation mechanisms, J. Econ. Theory 51 (1990) 233–254.
[11] P. Dasgupta, E. Maskin, Efficient auctions, Quart. J. Econ. 115 (2000) 341–388.
[12] J. Green, J.-J. Laffont, Characterization of satisfactory mechanisms for the revelation of preferences for public goods, Econometrica 45 (1977) 427–438.
[13] J. Green, J.-J. Laffont, Incentives in Public Decision Making, North-Holland, Amsterdam, 1979.
[14] T. Groves, Incentives in teams, Econometrica 41 (1973) 617–663.
[15] A. Heifetz, Z. Neeman, On the generic (im)possibility of full surplus extraction in mechanism design, Econometrica 74 (2006) 213–233.
[16] M. Jackson, A crash course in implementation theory, Soc. Choice Welfare 18 (2001) 655–708.
[17] P. Jehiel, M. Meyer-ter-Vehn, B. Moldovanu, W. Zame, The limits of ex-post implementation, Econometrica 74 (2006) 585–610.
[18] P. Jehiel, B. Moldovanu, Efficient design with interdependent valuations, Econometrica 69 (2001) 1237–1259.
[19] S. Johnson, J. Pratt, R. Zeckhauser, Efficiency despite mutually payoff-relevant private information: the finite case, Econometrica 58 (1990) 873–900.
[20] V. Krishna, Auction Theory, Academic Press, London, 2002.
[21] J. Marsden, M. Hoffman, Elementary Classical Analysis, W.H. Freeman and Co., New York, 1993.
[22] E. Maskin, Auctions and privatization, in: H. Siebert (Ed.), Privatization, Institut fur Weltwirtschaften der Universität Kiel, Kiel, 1992, pp. 115–136.
[23] H. Matsushima, Dominant strategy mechanisms with mutually payoff-relevant information and with public information, Econ. Letters 34 (1990) 109–112.
[24] H. Matsushima, Incentive compatible mechanisms with full transferability, J. Econ. Theory 54 (1991) 198–203.
[25] P. McAfee, P. Reny, Correlated information and mechanism design, Econometrica 60 (1992) 395–421.
[26] R. McLean, A. Postlewaite, Informational size and efficient auctions, Rev. Econ. Stud. 71 (2004) 809–827.
[27] A. Meirowitz, On the existence of equilibria to Bayesian games with non-finite type and action spaces, Econ. Letters 78 (2003) 213–218.
[28] C. Mezzetti, Mechanism design with interdependent valuations: efficiency, Econometrica 72 (2004) 1617–1626.
[29] Z. Neeman, The relevance of private information in mechanism design, J. Econ. Theory 117 (2004) 55–77.
[30] H. Royden, Real Analysis, third ed., Macmillan Publishing Company, New York, 1988.
[31] R. Selten, Axiomatic characterization of the quadratic scoring rule, Exper. Econ. 1 (1998) 43–61.
[32] W. Vickrey, Counterspeculation, auctions, and competitive sealed-tenders, J. Finance 16 (1961) 8–37.
[33] E. Zeidler, Nonlinear Functional Analysis and its Applications I: Fixed Point Theorems, Springer, New York, 1985.