Balanced-Budget Mechanisms with Incomplete Information[*]

Drew Fudenberg, David K. Levine, and Eric Maskin[**]

May 4, 1995

Abstract: We examine mechanism design with transferable utility and budget balance, using techniques we developed for the study of repeated games. We show that with independent types, budget balance does not limit the set of social choice functions that can be implemented. With correlated types and three or more players, budget balance is again not a constraint if no player has "too many" more possible types than any other player. Moreover, in the latter case, for generic probability distributions over types all social choice functions are implementable.

[*] We are grateful to Jacques Crémer for helpful comments, Pietro Veronesi for careful proofreading, and National Science Foundation Grants SBR 92-23320, 93-20695, and (Eric's) for financial support,
[**] Departments of Economics, Harvard, UCLA and Harvard.

1. Introduction

We study the problem of mechanism design under incomplete information when there is transferable utility. Using techniques originally developed in our study of repeated games, Fudenberg, Levine and Maskin [1994] (FLM) , we develop sufficient conditions for social choice functions to be Bayesian Nash implementable by mechanisms whose transfers sum to zero, that is, by balanced-budget mechanisms. We also consider when implementing mechanisms can be constructed that satisfy both budget-balance and interim individual-rationality constraints.

We have two sets of results. The first set concerns the case where the players' types are distributed independently. Here we show that *(i)* any social choice function (hereafter denoted *scf)* that can be implemented at all can be implemented with transfers that sum to zero, so that the balanced-budget requirement does not restrict the set of implementable scf's (see Proposition 1). Moreover, we show that *(ii)* if the scf being implemented yields higher *ex ante* expected social welfare than autarky, then the implementing balanced-budget mechanism can be constructed to be *ex ante* individually rational (see Proposition 3). Result *(i)* extends the theorem of d'Aspremont and Gérard-Varet [1979], and result *(ii)* generalizes a conclusion of Laffont and Maskin [1979]. Both those papers, unlike this one, consider only the first-best scf, and restrict attention to the "private values" case where each agent's utility function does not depend on the types of the others.[1]

Our second set of results concerns the case where the players' types may be correlated. Here we show that if there are three or more players, and no one player has "too many" more possible types than any other, then for generic prior probability distributions over types, *(iii)* any scf can be implemented with a balanced budget

[1] d'Aspremont and Gérard-Varet [1979] actually use a more general condition they call "condition C" instead of assuming independent types. However, verifying that the condition is satisfied requires solving a linear programming problem, and he only case in which they show the condition holds is that of indpependent types.

(Proposition 4) and *(iv)* any scf that increases *ex ante* social welfare can be implemented with a mechanism that is interim individually rational and has transfers whose expected value is nonpositive (Proposition 5). Result *(iii)* extends the theorems of Maskin [1986] and d'Aspremont, Crémer, and Gérard-Varet [1990], which restrict attention to private values and first-best allocations. In contemporaneous work, d'Aspremont, Crémer, and Gérard-Varet [1995] have obtained the same conclusion as our result (*iii)* on the weaker hypotheseses that there are at least three agents and each agent has at least two types. Result *(iv)* extends the results of Crémer and MacLean [1988] and McAfee, McMillan, and Reny [1989], who considered the special case of allocating a single private good, that is, an auction.

To prove our results, we use linear programming arguments to show that transfer rules satisfying both incentive compatibility and balanced-budget constraints can be found whenever it is possible to statistically distinguish between deviations by different players. We originally used these techniques to study repeated games, where the analog of utility transfers are the dynamic programming continuation payoffs. In the repeated game setting, balanced-budget conditions arise from the need to constrain continuation payoff vectors to lie in a particular hyperplane or half-space. Related techniques have been used in Legros [1988], Legros and Matsushima [1989], and Radner and Williams [1989] in the study of static moral hazard in teams.

2. The Model

There are I agents, $i = 1,\ldots, I$ with types $\theta_i \in \Theta_i$. We assume each type space Θ_i is finite, with n_i elements; we let $\Theta = \times_i \Theta_i$, with representative element $\theta = (\theta_1,\ldots,\theta_I)$, and set $n = \#\Theta = \prod_i n_i$. Players have a common prior probability distribution p on Θ.

An *outcome* is a pair $y = (x,t)$, where x is an allocation of physical goods (or, more generally, a public decision), and $t = (t_1,\ldots,t_I)$ is a vector of income transfers from the principal to each agent. We let X denote the (arbitrary) set of possible allocations, and let $T_i = \Re$ (the real line) be the set of feasible transfers to player *i*, so that $T \equiv \times_i T_i = \Re^I$. (The assumption of unbounded transfers is quite important for our results about correlated types; see the discussion at the end of section 5.)

Each agent $i = 1,\ldots,I$ has a quasilinear utility function $u_i(x,\theta) + t_i$. This formulation allows agent *i*'s utility to depend on the entire vector of types Θ; in the special case of *private values,* agent i's utility is not influenced by the types of other agents, so that utility is $u_i(x,\theta_i) + t_i$.

After observing his type, each agent sends a *message* or report m_i in the message space $\hat{\Theta}_i$ to an unmodelled principal; a *strategy* for player *i* is thus a map $m_i{:}\Theta_i \to \hat{\Theta}_i$. From the revelation principle, we restrict attention to direct revelation mechanisms, in which $\hat{\Theta}_i$ is isomorphic to the type space Θ_i; henceforth we identify these two spaces and write $\hat{\Theta}_i = \Theta_i$. Thus each agent *i* has $k_i = n_i^{\,n_i}$ pure strategies. We will find it convenient to number agent *i*'s types from 1 to n_i, and index strategies by vectors, so that $m_i{*} = m_i[1,2,\ldots.,n_i]$ is the "truthful" strategy, and the strategy $m_i[1,\ldots.,1]$ corresponds to player *i* always reporting that he is the first type.

Let *f* be a map from Θ to *X;* we interpret this map as a *social choice function* (scf). Let $F = X^{\Theta}$ denote the space of all such maps. A *mechanism* is a map $(f,\tau){:}\Theta \to X \times T$. Each mechanism induces a Bayesian game among the agents in the obvious way. If truthful reporting is an equilibrium of this game, we say that the mechanism *enforces* truthful reporting, and also that it *implements* the social choice function *f.* Formally, this requires that the incentive-compatibility constraint

(IC)
$$\sum_{\theta_{-i}} p(\theta_{-i}|\theta_i)[u_i(f(\theta_i,\theta_{-i}),(\theta_i,\theta_{-i})) + \tau_i(\theta_i,\theta_{-i})] \geq$$
$$\sum_{\theta_{-i}} p(\theta_{-i}|\theta_i)[u_i(f(\hat{\theta}_i,\theta_{-i}),(\theta_i,\theta_{-i})) + \tau_i(\hat{\theta}_i,\theta_{-i})]$$

hold for each player i, type θ_i and every possible report $\hat{\theta}_i$, where $p(\theta_{-i}|\theta_i)$ is the probability of θ_{-i} conditional on θ_i. We call this the *enforceability constraint* for the mechanism.[2] The mechanism (f,τ) has a *balanced budget* if the sum of the transfers equals zero for every possible vector of reports, that is $\sum_i \tau_i(\theta)=0$ for all profiles θ.

We will sometimes suppose that each player has the option of not participating in the mechanism, and further that the payoff from not participating does not depend on the realization of the types; in this case we normalize the non-participation value to 0. We distinguish two cases, depending on the point at which the players can opt not to participate. If the players can opt out after learning their types, the mechanism must be *interim individually rational:*

(Interim IR) for all players i and all θ_i

$$\sum_{\theta_{-i}} p(\theta_{-i}|\theta_i)[u_i(f(\theta_i,\theta_{-i}),(\theta_i,\theta_{-i}))+\tau_i(\theta_i,\theta_{-i})]\geq 0\,.$$

If the players must decide whether to participate before learning their type, the mechanism must satisfy the weaker condition of being *ex ante individually rational*:

(*Ex Ante* IR) for all i

$$\sum_{\theta_i} p(\theta_i)\sum_{\theta_{-i}} p(\theta_{-i}|\theta_i)[u_i(f(\theta_i,\theta_{-i}),(\theta_i\theta_{-i}))+\tau_i(\theta_i,\theta_{-i})]\geq 0\,.$$

An scf f is *extremal* if there are positive weights $\{\lambda_i\}$ such that f satisfies

(Extremal) $f\in\arg\max_{f'\in F}\sum_i \lambda_i \sum_\theta p(\theta)u_i(f'(\theta),\theta)$.[3]

---

[2]This is the "interim" version of the constraint, which corresponds to i's decision problem after he learns his own type. The "*ex ante*" version formulates the constraint on i's choice of strategy before he learns his type. These two formulations are equivalent when all types have positive probability.

[3]This condition is sometimes called *ex ante* Pareto optimality in the mechanism design literature. However, because utility is transferable, **(Extremal)** is not interpretable as an efficiency condition when the utility

We say that an scf is *ex ante socially rational* if the sum of the player's *ex ante* utilities is higher under the mechanism than under "autarky:"

(*Ex ante* SR) $\sum_{\theta}\sum_{i} p(\theta)u_i(f(\theta),\theta) \geq 0.$

This condition is clearly necessary in order for *f* to be implemented with a balanced-budget mechanism that is ex ante individually rational. Finally, an scf is *interim socially rational* if the social rationality constraint is imposed for all profiles of types:

(Interim SR) $\sum_{i} u_i(f(\theta),\theta) \geq 0 \quad \text{for all } \theta.$

3. Pairwise Identifiability

We begin by studying the probability distributions over reports induced by particular strategies. The key condition we use is called *pairwise identifiability.* Loosely speaking, a strategy profile is pairwise identifiable for players *i* and *j* if the probability distributions over reports resulting from deviations by player *i* can be statistically distinguished from those resulting from deviations by *j;* that is, if deviations by player *i* lead to distributions over outcomes that cannot be replicated by some (possibly random) reporting strategy of player *j*. Lemma 1 below shows that when this condition is satisfied, any implementable scf can be implemented with a balanced budget.

To state the condition precisely, choose a numbering scheme for each agent's pure strategies and fix a strategy profile $m = (m_1, \ldots, m_I)$ for the agents. We let $\Pi_i(m_{-i})$ be the $k_i \times n$ matrix whose *k*th row is the probability distribution over reports $\hat{\theta}$ when agent i uses his *k*th pure strategy and the other players use m_{-i}. For any pair of players *i* and *j*,

---

weights are not equal. See Ledyard and Palfrey [1994] for a discussion of the relationship between these two conditions under a different normalization of the utility functions.

we let $\Pi_{ij}(m)$ be the $(k_i + k_j) \times n$ matrix formed by stacking $\Pi_i(m_{-i})$ on top of $\Pi_j(m_{-j})$:

$$\Pi_{ij}(m) = \begin{bmatrix} \Pi_i(m_{-i}) \\ \Pi_j(m_{-j}) \end{bmatrix}.$$

Definition: The profile *m* is *pairwise identifiable for players i and j* if rank $(\Pi_{ij}(m))$ = rank$(\Pi_i(m_{-i}))$ + rank$(\Pi_j(m_{-j}))$ -1, where the rank of the matrix is the dimension of the space spanned by its rows.

Pairwise identifiability requires that the stacked matrix $\Pi_{ij}(m)$ have the largest rank possible given the ranks of $\Pi_i(m_{-i})$ and $\Pi_j(m_{-j})$: $\Pi_{ij}(m)$ cannot have rank equal to the sum of the ranks of the two submatrices, since the row of $\Pi_i(m_{-i})$ corresponding to *i* playing m_i must be the same as the row of $\Pi_j(m_{-j})$ corresponding to *j* playing m_j . (Note, though, that pairwise identifiability does not require that the constituent matrices themselves have full row rank. Indeed, submatrices *cannot* have full rank. As the proof of Lemma 3 demonstrates, the maximum possible rank of $\Pi_i(m_{-i})$ is $n_i(n_i - 1) + 1$).

Lemma 1: If the social choice function *f* is implementable by some mechanism, and the truthful reporting profile is pairwise identifiable for all pairs of agents i, j, then *f* can be implemented by a mechanism with a balanced budget.

Proof: This is essentially lemma 5.5 of FLM. That result shows that if a pure action profile (here, truthful reporting) is *enforceable*, and is pairwise identifiable for every pair of players, then it can be enforced with *continuation payoffs* (here, transfer payments) on any *regular hyperplane*. Enforceability in the current context is equivalent to the social choice function being implementable by some mechanism. Regular hyperplanes are defined by weights $\lambda_i \neq 0$, and the transfer payments lie on the regular hyperplane λ if $\sum_i \lambda_i t_i = 0$; a balanced budget corresponds to the particular regular hyperplane where all weights are equal.

❑

Here is a sketch of the argument behind Lemma 1. The first step is to note that a sufficient condition for f to be implementable with a balanced budget is that, for every pair of players *i,j* , and every pair of nonzero weights β_i, β_j, there be transfer functions τ_i, τ_j that enforce truthful reporting for i and j (that is, satisfy **IC** for these players**)** and also satisfy $\beta_i \tau_i(\theta) + \beta_j \tau_j(\theta) = 0$.[4] Setting all of the weights β_i to equal 1, and considering the pairs (1,2), (3,4) and so on, shows that this condition implies that f can be implemenated with budget balance for any even number of players. If there are an odd number of players, then define transfers for players (4,5), (6,7) and so on as above, with budget balance within each pair. To define the transfers for players 1,2, and 3, let $\{\tau_1, \tau'_2\}$ be transfers that enforce f for 1 and 2 and satisfy $\beta_1(\theta) + \frac{1}{2}\beta_2(\theta) = 0$, let $\{\tau''_2, \tau_3\}$ be transfers that enforce f for 2 and 3 and satisfy $\frac{1}{2}\beta_2(\theta) + \beta_3(\theta) = 0$, and set $\tau_2(\theta) = \frac{1}{2}\tau'_2(\theta) + \frac{1}{2}\tau''_2(\theta)$. Since τ_2 is a convex combination of transfer fuctions that enforces truthful reporting for player 2, τ_2 does also, and the transfers $\tau_1, \tau_2, \ldots, \tau_I$ satisfy budget balance, which finishes the verification of the first claim.

Next, note that the **(IC)** constraint for player k can equivalently be expressed as the inequalities

$$(*)\ \sum_{\theta} p(\theta)[u_k(f(\theta),\theta) + \tau_k(\theta)] \geq \sum_{\theta} p(\theta)[u_k(f(m_k(\theta_k),\theta_{-k}),\theta) + \tau_k(m_k(\theta_k),\theta_{-k})]$$

---

[4] FLM [1994] call this "enforceability with respect to pairwsie hyperplanes." The argument in this paragraph is essentially lemma 5.3 of that paper.

for all reporting strategies m_k

Consider the set of inequalities consisting of the incentive constraints for players i and j (as given by (*) when $k = i, j$). In light of the argument above, it will be sufficient to show that these constraints can be satisfied when we replace each τ_j in this set by $-\beta\tau_i$ for any nonzero β.

To investigate this question we turn to the matrices $\Pi_i(m^*_{-i})$ and $\Pi_j(m^*_{-j})$. Because $\Pi_i(m^*_{-i})$ does not have full rank, it has a row corresponding to some reporting strategy m_i' that can be written as a linear combination of the other rows. Because *f* is implementable, there exists a transfer rule τ_i that satisfies the incentive constraints (*) for the strategies corresponding to these other rows. Moreover, implementability implies that m_i' can be chosen so that τ_i also satisfies (*) for this strategy as well. Thus, in seeking a solution to the system of incentive constraints, we can delete the constraint corresponding to m_i'. A similar argument applies to $\Pi_j(m^*_{-j})$. Proceeding iteratively, we can delete enough rows from $\Pi_i(m^*_{-i})$ and $\Pi_j(m^*_{-j})$ that the reduced matrices have full rank. Pairwise identifiability then implies that the system of incentive constraints for i and j corresponding to the reduced versions of $\Pi_i(m^*_{-i})$ and $\Pi_j(m^*_{-j})$ is solvable. But the way that we have deleted rows ensures that a solution to this system satisfies *all* the incentive constraints for players i and j.

4. Independent Types

Suppose now that the prior distribution p on the agents' types is a product measure, $p(\theta) = \prod_i p_i(\theta_i)$, where $p_i(\theta_i)$ is the marginal probability of θ_i, so that the types of the various agents are independently distributed.

Lemma 2: With independent types, every pure strategy profile is pairwise identifiable for every pair of players.

Proof: The intuition for this result is that when types are independent, the joint distribution over report vectors is simply the product of the marginal distributions of the reports of the individual players, so that the report distributions arising from deviations by player *i* cannot be replicated by deviations of player j. More formally, with independent types, the game has a *product structure* in the sense of FLM. The concept of a product structure applies to a much broader class of games than those considered here. In a game in which each player *i* chooses an action a_i and the profile of actions $a=(a_i,\ldots,a_I)$ results in a distribution over outcome profiles $y=(y_1,\ldots,y_I)$, product structure is satisfied if (*i*) the marginal distribution of y_i given *a* depends only on a_i, and (*ii*) the joint distribution of *y* given *a* is the product of the marginal distributions. In our setting, an "action" is a reporting strategy and a profile of "outcomes" is just a profile of reports. Hence (*i*) is satisfied trivially (a player's report depends only on his own reporting strategy) and (*ii*) holds thanks to independence of types. The result then follows from Lemma 7.1 in FLM

❑

The following is an immediate consequence of Lemmas 1 and 2.

Proposition 1: If types are independently distributed, then any social choice rule that is implementable by some mechanism can be implemented by a mechanism with a balanced budget.

In light of Proposition 1, it is interesting to ask which social choice rules are implementable. The next proposition shows that in the case of private values *f* is implementable if it is extremal. Moreover, the set of extremal mechanisms does not depend on the particular distribution over types:

Proposition 2: *(a)* If social choice rule f is extremal under the measure p on Θ, it is extremal for all measures with the same support. If, moreover, *(b)* agents have private values, then any extremal f is implementable.

Proof: *(a)* If f is extremal, there are positive weights $\{\lambda_i\}$ such that f is a solution of $\max_{f'\in F} \sum_i \lambda_i \sum_\theta p(\theta) u_i(f'(\theta),\theta)$. Reversing the order of sums, we find that for each θ, $f(\theta)$ must maximize $\sum_i \lambda_i u_i(f'(\theta),\theta)$, that is, an extremal rule must maximize social utility pointwise.

(b) Suppose f is extremal and that agents have private values. For each player i, set $\tau_i(\hat{\theta}) = \sum_{j\neq i} (\lambda_j / \lambda_i) u_j(f(\hat{\theta}),\hat{\theta}_j)$, where the weights λ_j are from the proof of part *(a)*, and $\hat{\theta}$ is the observed vector of reports. If all others report truthfully, player i's overall payoff from reporting $\hat{\theta}_i$ when his type is θ_i is

$$u_i(f(\hat{\theta}_i,\hat{\theta}_{-i}),\theta_i) + \sum_{j\neq i} (\lambda_j / \lambda_i) u_j(f(\hat{\theta}_i,\hat{\theta}_{-i}),\hat{\theta}_j) = (1/\lambda_i) \sum_{j=1}^{I} \lambda_j u_j(f(\hat{\theta}_i,\theta_{-i}),\theta_j)$$

Consequently player i's objective is the same as the social maximization problem for which f is extremal; hence player i's payoff is maximized by reporting truthfully, so that for each θ the outcome is that prescribed by f.

❑

Remark 1: The proof of part (b) shows that, with private values, any extremal mechanism is implementable in dominant strategies. When the utility weights $\{\lambda_i\}$ on all players are equal, the transfers constructed in the proof correspond to a Groves mechanism (Groves [1973].) Propositions 1 and 2 together imply that any extremal mechanism is implementable by a mechanism with a balanced budget in the case of private values and independent types; this result was first obtained by d'Aspremont and Gérard-Varet [1979].[5]

Remark 2: Depending on the utility functions, many non-extremal social choice functions may be implementable. One classic example is that of monotonic scfs: suppose that X and each Θ_i are subsets of the real line, that each u_i is concave and twice-differentiable in *x,* and that $\partial u_i / \partial x$ is increasing in θ_i. Then any scf that is differentiable and nondecreasing in θ_i for each i is implementable. (Laffont and Maskin [1979]). Conversely, if $\partial u_i / \partial x$ is increasing in θ_i, but there are not private values, then some extremal scf's may not have the "expected monotonicity" condition (namely that $E_{\theta_{-i}} f(\theta_i, \theta_{-i})$ be nondecreasing in θ_i) and consequently may not be implementable when types are independent.[6] Section 5 gives a sufficient condition for *any* scf to be implementable in the case of correlated types.

Next we turn to the question of implementation under individual rationality constraints.

Proposition 3: If types are independently distributed, and the social choice function f is *ex ante* socially rational and implementable, it is implementable by a balanced-budget mechanism that is *ex ante* individually rational.

[5] As we noted in the introduction, their paper shows that the assumption of independence can be relaxed to "condition C" on the solution to a linear program, but the paper does not indicate what conditions other than independence might suffice for condition C.

[6] As an example, suppose that $I = 2$ and $u_i(x, \theta) = (\theta_i - 2\theta_j)x$, where x can be either 0 or 1 and each θ_i takes values on the grid $\{-1, -1 + 1/k, \ldots, 1 - 1/k, 1\}$. The scf that chooses x to maximize the sum of the two player's utilities has x = 1 only if the sum of the two types is negative, and consequently is not monotonic. It is straightforward to check that this scf is consequently not implementable if types are independently distributed and the "grid" of types is sufficiently fine, (i.e., k is sufficiently big).

Remark: This proposition generalizes a similar result of Laffont and Maskin [1979], who assumed private values.

Proof: Under the hypotheses, Proposition 1 implies that there exists balanced transfer rule τ such that (f,τ) implements f. Because transfers are balanced and f is *ex ante* socially rational, we have

$$\sum_i \sum_\theta p(\theta)[u_i(f(\theta),\theta)+\tau_i(\theta)] \geq 0.$$

That is, the sum over players of the "direct" expected utility plus the expected transfer is positive. Thus we can find constants $\{k_i\}$ independent of θ such that $\sum_i k_i = 0$, and $\sum_\theta p(\theta)[u_i(f(\theta),\theta)+\tau_i(\theta)+k_i] \geq 0$ for all i. Hence the transfers $\tau_i'(\theta)=\tau_i(\theta)+k_i$ implement f, sum to 0, and are *ex ante* individually rational.

❑

Remark: The results of Laffont and Maskin [1979] and Myerson and Satterthwaite [1981] show that Proposition 3 does not hold if we replace *ex ante* individual rationality with the stronger requirement of interim IR, even when the scf to be implemented is interim socially rational. Indeed, these results establish that even if the balanced budget requirement is weakened so that only the condition that the expected sum of the transfers be nonpositive is required, there is still a conflict with interim IR. (However, Proposition 5 shows that this conflict vanishes in the case of correlated types.)

5. Correlated Types

Now we allow the prior distribution p to be any probability distribution on Θ. Here we can show that for generic distributions, truthful reporting is pairwise identifiable and every social choice function is implementable. We prove the following result in the Appendix:

Lemma 3: Suppose that

(♣) $\prod_{k \neq i,j} n_k \geq n_i$ for all $i \neq j$.

Then for generic probability distributions on Θ, the truthful reporting profile is pairwise identifiable for every pair of players.

Remark: An example may help to show both why the lemma requires more than two players and why pairwise identifiability need not be satisfied by all strategy profiles. Suppose that there are two players, each with two types. Then, for any m_2, the rows corresponding to the three strategies $m_1[1,1]$, $m_1[2,2]$, and $m_1[1,2]$ are generically a basis for the matrix $\Pi_1(m_2)$, (i.e., the matrix generically has rank three). Similarly for $\Pi_2(m_1)$. (This claim is verified in the proof in the Appendix.) Then pairwise identifiability requires that $\Pi_{12}(m)$ have rank 5. But there are only 4 possible report profiles -- that is, $\Pi_{12}(m)$ has only 4 columns and so rank 5 is impossible. Hence, there are deviations by player 1 that cannot be distinguished from those by player 2. However, if we introduce a third player who reports truthfully, and whose types are correlated with those of player 1 differently from their correlation with those of player 2 (the generic case), then we may be able to use the distribution over messages by player 3 to distinguish between player 1's and 2's deviations. Note, however, that if the third player does not tell the truth, but instead reports type 1, say, regardless of his true type, then his report contains no information about the strategies of other players.[7] Thus, even with three players, pairwise identifiability for players 1 and 2 is not satisfied by all profiles.

[7] More generally, whenever player 3 makes the same announcement for two different types, there is a loss of information about the players 1 and 2.

If each player has the same number of possible types, than condition (♣) *is* satisfied whenever there are three or more players. More generally, the condition is satisfied if the number of players is sufficiently large compared to the variation in the number of types across players; a sufficient condition along these lines is that the number of players be at least $2 + \ln(\max_i n_i)/\ln(\min_i n_i)$.

Once we have established that the truthful reporting strategies are pairwise identifiable, we know that any implementable social choice function is implementable with a balanced budget. But, as the next lemma establishes, any social choice function is generically implementable provided that a weak condition on the number of types holds.

Lemma 4: Suppose that, for all i, $\prod_{j \neq i} n_j \geq n_i - 1$. Then for generic probability distributions on Θ, any social choice function f is implementable.

Remark: Note that the generic set guaranteed by the lemma is independent of the social choice rule specified. The proof shows that each fixed f is implementable for a generic distribution; since there are only finitely many social choice functions, and a finite intersection of open and dense sets is open and dense, the stated conclusion follows. Note also that the hypothesis of lemma 4 is implied by hypothesis (♣) of lemma 3.

Informal proof: To show that f is implementable by some mechanism, it suffices to consider each player i separately and show that there are transfers that make it optimal for i to report truthfully if he believes that his opponents will do so. If player i is the only player, the conclusion of the lemma is easily seen to be false: suppose, for example, that player i prefers allocation x to allocation y by an amount equal to θ_i, where $\theta_i = 1, 2, 3$, and that the social choice function calls for allocation x if $\theta_i = 1$ and y otherwise. Then any transfer that makes type $\theta_i = 3$ willing to report truthfully will induce all types to make a report that leads to *y* being chosen. Even if there are other players, we are no better off if their types are independent, since in that case, their reports embody no useful information about i's type.[8]

The principle behind lemma 4 is that correlation between types of the various players allows the design of mechanisms that "detect and punish" misreporting by player *i*. More precisely, the joint distribution of others' reports (assuming that they report truthfully) is different depending on player i's type. Hence i can be induced to report truthfully by making his transfer depend appropriately on these other reports. To accomplish this, each type $\theta_i' \neq \theta_i$ must be deterred from reporting θ_i. That is, each type θ_i' must give rise to a different conditional distribution on others' reports. Since there are $n_i - 1$ such types we need to deter, and $\Pi_{j\neq i} n_j$ types of other players, we are led to the hypothesis used in lemma 4.

Proof: We must show that there is a transfer function τ that satisfies (IC) for all i and all $\hat{\theta}_i \neq \theta_i$. That is, we must satisfy $n_i(n_i - 1)$ constraints. Now these constraints can be satisfied provided that the matrix $\Pi_i(m^*_{-i})$ has rank at least $n_i(n_i - 1)$ (where m^*_{-i} is the profile of truthful reporting strategies for players other than i). But, from the proof of Lemma 3, $\Pi_i(m^*_{-i})$ generically has rank at least $n_i(n_i - 1)$ provided that $n \geq n_i(n_i - 1)$, that is, $\Pi_{j\neq i} n_j \geq n_i - 1$.

❑

Combining lemmas 3 and 4 yields the following result.

Proposition 4: Suppose that (♣) holds. Then for generic probability distributions on Θ, any social choice function is implementabile by a balanced-budget mechanism.

---

[8] In this sort of setting only monotonic social choice functions can be implemented, see Remark 2, following Proposition 2.

Proposition 4 generalizes the conclusions of d'Aspremont, Crémer, and Gérard-Varet [1990] and Maskin [1986], who assumed private values and extremal social choice functions. Maskin assumed each player had only 2 types; d'Aspremont, Crémer and Gérard-Varet studied generic probability distributions, as we do. Using a technique based on "scoring rules" instead of linear algebra, d'Aspremont, Crémer, and Gérard-Varet [1995] obtain the same conclusion as Proposition 4, replacing our hypothesis (♣) with the weaker condition that there are at least three players, each of whom has at least two types.

Next we turn to the question of implementation under individual rationality constraints. Proposition 3 readily extends to the case of correlated types. That is, under the hypotheses of proposition 4, any *ex ante* SR social choice function can be implemented for generic distributions on types by a balanced-budget mechanism that satisfies ex ante individual rationality. However, as in the case of independent types, one can construct examples in which generically implementation with a balanced budget and interim individual rationality is impossible, even if the scf satisfies interim SR; see the example following Proposition 5.

Nevertheless, unlike the case of independent types, interim individual rationality is attainable if the balanced-budget requirement is weakened to require only that the sum of the transfers have a non-positive expected value. This weaker condition is of interest if, for example, the agents have access to a risk-neutral bank that can provide an infusion of funds in some states of the world, but will agree to participate in the mechanism only if its expected profit is non-negative.

Proposition 5: Suppose that $\prod_{j \neq i} n_j \geq n_i$ for all i. Let f be a social choice function that satisfies *ex ante* social rationality. Then for a generic probability distribution over θ, f is implementable with a mechanism that is interim individually rational and such that the expected value of the sum of the transfers is nonpositive.

Remark: The proof shows that *f* can be implemented with transfers such that every player's interim individual rationality constraint holds with exact equality for each of the player's possible types; the assumption that the social choice rule is *ex ante SR* then implies that the expected sum of the transfers is nonnegative. Thus, if we think of the negative of the transfers as accruing to an unmodelled "bank" or principal, this principal extracts *all* the interim surplus that the social choice rule provides to the players. When the types of different players are independently distributed, mechanisms that enforce truth telling may need to leave some surplus ("informational rents") to those types with a relatively high value from participating in the mechanism. This is why proposition 5 is false for the case of independent types, as shown by the results of Laffont and Maskin [1979] and Myerson and Satterthwaite [1981].

Proof: Fix a player i . *A*s the Remark indicates, it suffices to show that there exists a transfer rule τ_i that satisfies (**IC**) and (**interim IR**) with equality for all values of θ_i. For each $k = 1,\ldots,n_i$, let ρ_k be the n-dimensional vector of coefficients corresponding to the interim IR constraint for the *k*th value of θ_i. **(Interim IR)** can be expressed as

$$\sum_{\theta_{-i}} p(k,\theta_{-i})\tau_i(k,\theta_{-i}) \geq -\sum_{\theta_{-i}} p(k,\theta_{-i})u_i(f(k,\theta_{-i}),(k,\theta_{-i})),$$

where we are slightly abusing notation by using "k" to denote the *k*th value of θ_i. Hence ρ_k, the vector of coefficients of the transfers, is given by

$$\begin{array}{c} (0,\ldots,0,p(k,\cdot),0,\ldots,0), \\ \uparrow \\ k\text{th entry} \end{array}$$

where each "0" corresponds to a vector of $\prod_{j \neq i} n_j$ zeros). Similarly, let ρ_{kl} be the vector of coefficients corresponding to the IC

$$\sum_{\theta_{-i}} p(k,\theta_{-i})\tau_i(k,\theta_{-i}) + \sum_{\theta_{-i}} (-p(k,\theta_{-i})\tau_i(l,\theta_{-i}))$$

$$\geq \quad \sum_{\theta_{-i}} p(k,\theta_{-i})(u_i(f(l,\theta_{-i}),(k,\theta_{-i})) - u_i(f(k,\theta_{-i}),(k,\theta_{-i})))$$

Hence, ρ_{kl} can be expressed as

$$(0,\ldots,0, \underset{k\text{th entry}}{\underset{\uparrow}{p(k,\cdot)}} ,0,\ldots,0, -\underset{l\text{th entry}}{\underset{\uparrow}{p(k,\cdot)}} ,0,\ldots,0).$$

To show that the set of these vectors is linearly independent, we must demonstrate that if the equation

$$(**) \qquad \sum_{k=1}^{n_i} \lambda_k \rho_k + \sum_{k=1}^{n_i} \sum_{l \neq k} \lambda_{kl} \rho_{kl} = 0$$

holds, then all the scalars λ_k are λ_{kl} are zero.

Fix θ_{-i} and consider the component of (**) corresponding to the hth value of θ_i (i.e., the coefficient of the transfer $\tau_i(h,\theta_{-i})$. We have

$$\left(\sum_{l \neq h} \lambda_{hl} + \lambda_h\right) p(h,\theta_{-i}) + \sum_{k \neq h} (-\lambda_{kh}) p(k,\theta_{-i}) = 0.$$

From the inequality in the statement of the Proposition, we know that the vectors $\{p(1,\cdot),\ldots,p(n_i,\cdot)\}$ are generically, linearly independent. Hence, since the above equation holds for all θ_{-i}, we have

$$\sum_{l\neq k}\lambda_{hl}+\lambda_h=-\lambda_{kh}=0 \text{ for } k\neq h.$$

We conclude that all the scalars are indeed zero.

❑

This result generalizes those of Crémer and McLean [1988] and McAfee and Reny [1988] for the case of auctions with private values.

Discussion: As with previous results about correlated types, propositions 4 and 5 cannot be strengthened to impose a uniform bound on the absolute values of the transfers; the transfers required can grow without bound along a sequence of correlated distributions whose limit is a distribution with independent types. For this reason the results are of the most interest in cases where the amount of correlation is not negligible.

To see that Proposition 5 cannot be strengthened to require an exactly balanced budget, consider the following example.

Example: Suppose that there are two players $(i=1,2)$, each with two possible types, θ_i^1 and θ_i^2. For each $i=1,2$, let

$$u_i\left(f\left(\hat{\theta}_1,\hat{\theta}_2\right),\theta_i\right)=\begin{cases}0, \text{ if } \theta_i=\hat{\theta}_i\\ 1, \text{ if } \theta_i\neq\hat{\theta}_i\end{cases}$$

Notice that f is not only *ex ante* but interim SR. Moreover, the example satisfies the hypotheses of Proposition 5. Nevertheless, for an open set of priors f cannot be implemented by a balanced budget mechanism let alone one satisfying the interim IR constraints.

To see this, take

$$\begin{pmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{pmatrix} = \begin{pmatrix} \frac{1}{3} & \frac{1}{6} \\ \frac{1}{6} & \frac{1}{3} \end{pmatrix} ,$$

where, for all i and j, $p_{ij} = p_r\{\theta_1^i, \theta_2^j\}$. If the budget balances, we can express the incentive-compatibility constraints for Player 1 and 2 as

$$\begin{aligned} \frac{2}{3}t_{11} + \frac{1}{3}t_{12} - \frac{2}{3}t_{21} - \frac{1}{3}t_{22} &\geq 1 \\ -\frac{1}{3}t_{11} - \frac{2}{3}t_{12} + \frac{1}{3}t_{21} + \frac{2}{3}t_{22} &\geq 1 \\ -\frac{2}{3}t_{11} + \frac{2}{3}t_{12} - \frac{1}{3}t_{21} + \frac{1}{3}t_{22} &\geq 1 \\ \frac{1}{3}t_{11} - \frac{1}{3}t_{12} + \frac{2}{3}t_{21} - \frac{2}{3}t_{22} &\geq 1 \end{aligned}$$

Adding the inequalities in (*) together, we obtain $0 \geq 4$, a contradiction. Clearly, the same contradiction would hold if we perturbed the probabilities slightly.

Appendix

.

Recall that the definition of pairwise identifiability for players i and j compares the rank of a stacked matrix to those of its two constituent submatrices, and requires that $\mathrm{rank}(\Pi_{ij}(m)) = \mathrm{rank}(\Pi_i(m_{-i}) + \mathrm{rank}(\Pi_j(m_{-j})) - 1$. That is, the subspace corresponding to the intersection of the spans of the two matrices has dimension 1. (The span of a matrix is the linear space comprising all linear combinations of its rows.)

Lemma 3: Suppose that for all i and j, $i \neq j$,

(♣) $\prod_{k \neq i,j} n_k \geq n_i$

Then for generic probability distributions on Θ, the truthful reporting profile m^* is pairwise identifiable for every pair of players.

Proof: Fix the truthful reporting profile m^*. We must show that

$$\dim\left[\mathrm{span}(\Pi_i(m^*_{-i})) \cap \mathrm{span}(\Pi_j(m^*_{-j}))\right] = 1 \quad .$$

Part 1: Define the strategies $\overline{m}_i(k) = m_i[k, k, \ldots, k]$ (player i makes the same report k regardless of his type) and, for all $h, k \neq 1$, $\overline{m}_i(h,k) = m_i[\underbrace{1,\ldots,1,\mathrm{k},1}_{\text{h-1 components}},\ldots,1]$ (player i reports type 1 unless he is type $h \neq 1$, in which case he reports type $k \neq 1$.) Let $\overline{M}_i$ consist of all strategies $\overline{m}_i(k)$ and $\overline{m}_i(h,k)$. For each m_i, let $\rho(m_i)$ denote the corresponding row of $\Pi_i(m^*_{-i})$ (that is, the distribution over vectors of reports generated by the profile (m_i, m^*_{-i})). Then

$$\rho(m_i[k_1, k_2, \ldots, k_{n_i}]) = \rho(\overline{m}_i(k_1)) + \sum_{j=2}^{n_i} \left(\rho(\overline{m}_i(j, k_j)) - \rho(\overline{m}_i(j, k_1))\right).$$

This shows that the span of $\rho(\overline{M}_i)$ $(= \bigcup_{\overline{m}_i \in \overline{M}_i} \rho(\overline{m}_i))$ and $\Pi_i(m*_{-i})$ are the same.[9] In particular, the rank of $\Pi_i(m*_{-i})$ can be at most the number of elements of $\overline{M}_i$, which is $(n_i - 1)^2 + n_i$. (In the example following the statement of Lemma 3 in the text, each player has two types, so that the rank of $\Pi_i(m*_{-i})$ can be at most 3.)

Part 2: We will find it convenient to work with the projection operator that sends probability distributions on Θ to the corresponding marginal distributions on the space of type vectors θ_{-i}. This map, denoted H_i, is given by the $n \times \prod_{j \neq i} n_j$ matrix

$$H_i(\theta, \theta'_{-i}) = \begin{cases} 1 \text{ if } \theta_{-i} = \theta'_{-i} \\ 0 \text{ if } \theta_{-i} \neq \theta'_{-i} \end{cases}.$$

If p is a probability distribution on types, that is, an $n-$dimensional probability vector, then $\sum_{\theta} p(\theta) H_i(\theta, \cdot)$ is the corresponding marginal distribution on types of players other than player i. Denote this marginal distribution by p_{-i}.

Claim: For generic p, all i, and all $j \neq i$, rank $\Pi_i(m*_{-i}) = \text{rank}[\Pi_i(m*_{-i})H_j] = \#\overline{M}_i$.

To establish the claim, we will show that for generic p, the row vectors in $\rho(\overline{M}_i)H_j$ are linearly independent. This will imply that the rows of $\rho(\overline{M}_i)$ are independent as well; the claim then follows from the conclusion of part 1, that the spans of $\overline{M}_i$ and $\Pi_i(m^*_{-i})$ are equal.

Suppose to the contrary that there is a linear dependence, so that for some vector $\lambda \neq 0$, $\lambda\rho(\overline{M}_i)H_j = 0$. Let $p_{-j}(h, \theta_{-i-j})$ be the probability (under distribution p_{-j}) that player *i* is of type *h*, and that the types of players other than *i* and *j* are given by θ_{-i-j}.

[9]This conclusion holds for any profile m_{-i}, and not just for the truthful reporting profile.

Since in $\Pi_i(m^*_{-i})$ all players except i are telling the truth the rows of $\rho(\overline{M}_i)H_j$ are given by

$$
\begin{aligned}
\left(\rho(\overline{m}_i(k))H_j\right)(k,\theta_{-i-j}) &= \sum_{h=1}^{n_i} p_{-j}(h,\theta_{-i-j}), k=1,\ldots,n_i \\
\left(\rho(\overline{m}_i(\ell))H_j\right)(k,\theta_{-i-j}) &= 0, \qquad k\neq \ell, \quad k,\ell=1,\ldots,n_i \\
\left(\rho(\overline{m}_i(h,k))H_j\right)(k,\theta_{-i-j}) &= p_{-j}(h,\theta_{-i-j}), h,k\neq 1 \\
\left(\rho(\overline{m}_i(h,\ell))H_j\right)(k,\theta_{-i-j}) &= 0, \qquad h,\ell\neq 1, \quad k\neq \ell,1
\end{aligned}
$$

where $\rho(m_i)H_j(k,\theta_{-i-j})$ is the component of $\rho(m_i)H_j$ corresponding to report profile $(k,\theta_{-i,-j})$. Consequently, the component of $\lambda\rho(\overline{M}_i)H_j$ corresponding to (k,θ_{-i-j}) is

$$
\lambda(k)\sum_{h=1}^{n_i} p_{-j}(h,\theta_{-i-j}) + \sum_{h=2}^{n_i} \lambda(h,k)p_{-j}(h,\theta_{-i-j}) = 0
$$

where we are using the fact that the strategies $\overline{m}_i(1,k)$ are not in $\overline{M}_i$. This may be rewritten as

$$
\lambda(k)p_{-j}(1,\cdot) + \sum_{h=2}^{n_i}(\lambda(k)+\lambda(h,k))p_{-j}(h,\cdot) = 0.
$$

It follows that $\lambda = 0$ provided that the vectors $\left\{p_{-j}(1,\cdot),\ldots,p_{-j}(n_i,\cdot)\right\}$ are inequalities. Now the vectors $\left\{p_{-j}(1,\cdot),\ldots,p_{-j}(n_i,\cdot)\right\}$ form a $n_i \times \prod_{k\neq i,j} n_k$ matrix, and, given inequalities (♣) , such a matrix generically has rank n_i, that is, has independent rows as required.[10] This does not yet establish the claim, since it hypothesized generic probabilities on the set Θ, as opposed to generic marginal probabilities..

[10] Note that the requirement that probabilities sum to 1 plays no role in the arguments that follow, as every matrix is a scalar multiple of a matrix satisfying the adding up restriction, and scalar multiplication does not effect the rank of a matrix. The positivity requirement of probabilities similarly plays no role, as the set of

However, we may view the probabilities $p(\theta)$ as a $n_i \times \prod_{k \neq i} n_k$ matrix, in which rows correspond to types of player i , and columns to profiles of other player types. Let H_{ij} be the $\prod_{k \neq i} n_k \times \prod_{k \neq i,j} n_k$ projection matrix that computes marginal probabilities over types of players other than players i and j. That is, H_{ij} is the analog of H_i for maps from Θ_{-i} to Θ_{-i-j}. It is clear that that H_{ij} is surjective (onto). Consequently, since it is a linear map, it is an open map, and therefore preserves genericity. The observation that $p_{-j}(\theta_{-j}) = p(\theta_{-j}, \cdot)H_{ij}$ then completes the proof of the claim.

Part 3: The claim of part 2 shows that for generic p, $\dim[\text{span}\Pi_i(m^*_{-i})] = \dim[\text{span}(\Pi_i(m^*_{-i})H_j)]$, so that H_j is 1 to 1 on the span of $\Pi_i(m^*_{-i})$. Consequently, H_j is 1-1 on $\text{span}\Pi_i(m_{-i}^*) \cap \text{span}\Pi_j(m_{-j}^*)$, and so

$$\dim[\text{span}\Pi_i(m_{-i}^*) \cap \text{span}\Pi_j(m^*_{-j})] = \dim[\left((\text{span}\Pi_i(m^*_{-i}) \cap \text{span}\Pi_j(m^*_{-j}))\right)H_j].$$

Moreover,

$$\begin{aligned}
&[\text{span}\Pi_i(m^*_{-i}) \cap \text{span}\Pi_j(m^*_{-j}))]H_j \subseteq \\
&(\text{span}\Pi_i(m^*_{-i}))H_j \cap (\text{span}\Pi_j(m^*_{-j}))H_j \subseteq \\
&(\text{span}\Pi_j(m^*_{-j}))H_j
\end{aligned}$$

Now $\dim \text{span}\Pi_j(m^*_{-j})H_j = 1$. This follows since player j's strategy is irrelevant for the distribution of outcomes of types θ_{-j} : every strategy for player j leads to the same distribution over θ_{-j} as m^*_{-j} itself. Since $\dim(\text{span}\Pi_i(m^*_{-i}) \cap \text{span}\Pi_j(m^*_{-j})) \geq 1$, we conclude that $\dim(\text{span}\Pi_i(m^*_{-i}) \cap \text{span}\Pi_j(m^*_{-j})) = 1$.

❏

---

matrices with non-negative entries is the union of an open set with its boundaries, and consequently inherits properties that are generic in the space of all matrices.

References

Crémer, J. and R. McLean [1988] "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," Econometrica 56, 1247-1258.

d'Aspremont, C., J. Crémer, and L.-A. Gérard-Varet [1990] "Incentives and the Existence of Pareto-Optimal Revelation Mechanisms," Journal of Economic Theory, 51, 233-254.

d'Aspremont, C., J. Crémer, and L.-A. Gérard-Varet [1995] "Correlation, Independence, and Bayesian Incentives," mimeo.

d'Aspremont, C. and L.-A. Gérard-Varet [1979] "Incentives and Incomplete Information," Journal of Public Economics, 11, 25-45.

Fudenberg, D. and D. K. Levine [1988] "The Folk Theorem in Repeated Games with Unobservable Actions,", mimeo.

Fudenberg, D., D.K. Levine, and E. Maskin [1994] "The Folk Theorem with Imperfect Public Information," Econometrica, 62, 997-1039.

Fudenberg, D. and E. Maskin [1986] "Discounted Repeated Games with One-Sided Moral Hazard," mimeo.

Laffont, J.-J. [1989], Fundamentals of Public Economics, MIT Press.

Ledyard, J. and T. Palfrey [1994] "Voting and Lottery Drafts as Efficient Public Mechanisms," Review of Economic Studies, 61, 3327-356.

Laffont, J.-J. and E. Maskin [1979] "A Differentiable Approach to Dominant Strategy Mechanisms," in J.-J. Laffont, ed., Aggregation and the Revelation of Preferences, North-Holland.

Legros, P. [1988] "Sustainability in Partnerships," mimeo, Caltech.

Legros, P. and H. Matsushima [1991] "Efficiency in Partnerships," Journal of Economic Theory, 55,296-320.

Maskin, E. [1986], "Optimal Bayesian Mechanisms," in W. Heller, R. Starr, and D. Starrett, eds., Uncertainty, Information, and Communication, Cambridge University Press.

McAfee, R., J. McMillan, and P. Reny [1988] "Extracting the Surplus in a Common-Value Auction," Econonometrica 57 1451-1459.

McAfee, R., and P. Reny [1988] "Correlated Information and Mechanism Design," 60, 395-421.

Myerson, R. and M. Satterthwaite [1983] "Efficient Mechanisms for Bilateral Trading," Journal of Economic Theory, 28, 265-281.

Williams, S. and R. Radner [1988] "Efficiency in Partnership when the Joint Output is Uncertain," mimeo.