Efficient Design with Multidimensional, Continuous Types, and Interdependent Valuations

Scott Johnson, Nolan H. Miller, John W. Pratt, Richard J. Zeckhauser.*

January 26, 2004†

Abstract

We consider mechanism design in social choice problems in which agents' types are continuous, multidimensional, and mutually payoff-relevant, and there are three or more agents. If the center receives a signal that is stochastically related to the agents' types and direct returns are bounded, then for any decision rule there is a balanced transfer function that ensures that any strategy that is not arbitrarily close to truthful is dominated by one that is. If direct returns are continuous as well, truthful revelation becomes an ε-dominant strategy, all Bayes-Nash equilibrium strategies are nearly truthful, and at least one such strategy exists. If the center's information is not informative but agents' types are stochastically related, then there exist balanced transfers under which truthful revelation is a Bayesian ε-equilibrium, again for any decision rule. Analogous results hold for decentralized decision problems when agents also take mutually payoff-relevant actions in advance of any action by the center.

Keywords: Mechanism Design, Interdependent Values, Multidimensional Types, Efficient Mechanisms.

JEL Classification: C70, D60, D70, D82.

---

*Johnson (Australian National University, deceased); Miller (Kennedy School, Harvard University); Pratt (Harvard Business School); and Zeckhauser (Kennedy School, Harvard University). Please direct correspondence to Nolan Miller at 79 JFK St, Cambridge, MA, 02138 or nolan_miller@harvard.edu.

†A previous version of this paper circulated under the title "Efficient Design with Interdependent Valuations and an Informed Center." We thank Drew Fudenberg, Jerry Green, John H. Lindsey II, Rich McLean, Zvika Neeman, Andy Postlewaite, Phil Reny, Bill Sandholm, and Michael Schwarz for helpful comments.

1 Introduction

Eliciting private information about preferences to guide social decisions is a classic problem of economic theory. For the private-values case, in which agents' private information does not affect other agents' preferences, the pioneering work of Vickrey (1961), Clarke (1971), and Groves (1973) shows that if each agent's preferences depend only on his own information and if the budget need not be balanced, externality payments make honest revelation a dominant strategy. However, dominant strategy implementation is generally incompatible with the requirement that the budget balance.[1] If the solution concept is weakened, positive results are possible. For example, d'Aspremont and Gérard-Varet (1979, 1982) show in the private-values environment that if the agents' beliefs about other agents' types satisfies 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, Crémer, and Gérard-Varet (1990) 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.

When one agent's private information affects other agents' preferences, the case of interdependent valuations or mutually payoff-relevant private information, the problem becomes more difficult, and positive results have been mostly limited to the case where agents' types take on only finitely many values.[3] For example, Aoyagi (1998) considers a model with finite types and interdependent values, and he 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 agents' types are continuous and values are interdependent, some positive implementation results are possible, but only at the expense of imposing additional structure on the model. For example, in the auctions context, Maskin (1992) and Dasgupta and Maskin (2000) show that when buyers' values are interdependent and one dimensional, a generalized Vickrey auction is efficient if a single-crossing property (each buyer's signal to has a greater impact on his own value than on others'

[1] See Green and Laffont (1977; 1979) and the discussion in Mas-Colell, Whinston, and Green (1995).

[2] Unlike its use in implementation theory (see Jackson, 2001), 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. When we wish to say that all outcomes of the game agree with the decision rule, we will use the phrase "uniquely implements."

[3] Work in this area includes Crémer and McLean (1985; 1988); Johnson, Pratt, and Zeckhauser (1990); Matsushima (1990; 1991); and McLean and Postlewaite (2003).

values) is satisfied. When types are multidimensional and continuous, the problem becomes even more difficult because, in general, single-crossing will not be satisfied. In this case, Maskin (1992) and Dasgupta and Maskin (2000) argue that when buyers' types are "truly" multidimensional in the sense that they cannot be summarized by a one-dimensional type, there may be no efficient auction.

In a general mechanism design framework, Jehiel and Moldovanu (2001) explore the difficulty in implementing efficient decision rules when types are multidimensional and continuous. When agents' types are independently distributed, efficient design is possible only when a certain "congruence condition relating the social and private rates of information substitution is satisfied (Jehiel and Moldovanu, 2001, p. 1237)." Intuitively, this is because, generically, the center's one-dimensional transfers are not sufficiently rich to extract agents' private information when multiple dimensions of the agents' signals are important in determining their valuations.[4]

The present paper addresses the social choice problem in environments in which agents' private information is continuous, multidimensional, and mutually payoff-relevant (i.e., valuations are interdependent). However, we relax the Jehiel-Moldovanu assumption that agents' private information is independently distributed. Our primary concern is to show that when there are more than two agents and the center observes information that is stochastically dependent on agents' types (either because the center's signal is dependent on agents' types or because agents' types are, themselves, stochastically dependent), it is possible to design a system of transfer payments that induces agents to (nearly) truthfully reveal their private information and that (nearly) implements any decision rule. Thus our results provide a type of converse to those of Jehiel and Moldovanu. When agents' types are dependent – in a weak sense of dependence we call stochastic relevance – then (nearly) efficient design generally is possible.[5] Further, if the distribution of types satisfies stochastic relevance, our implementation results place very few additional requirements on agents' preferences.[6] In particular, we do not require a single-crossing property.

McAfee and Reny (1992) also consider the case of continuous, multidimensional, and mutually payoff-relevant types with stochastically dependent information. Taking the game played by the

[4] See Maskin (1992), Dasgupta and Maskin (2002), and Krishna (2002, section 17.2) for further discussion of this point in the auctions context.

[5] Jehiel and Moldovanu (2001) note that, despite their impossibility results, successful design may be possible if agents' types are correlated.

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

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.[7] 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 (Maskin (1992), Dasgupta and Maskin (2000), Krishna (2002)). Therefore, in such environments, the McAfee-Reny mechanism would be unable to extract the full information rent (i.e., the rents that would be generated if the auctioneer knew the agents' types), since the McAfee-Reny 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, thus making it possible to apply the McAfee-Reny result.

McAfee and Reny's (1992) positive results as well as those of Crémer and McLean (1985; 1988) 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 whenever learning an agent's type provides information about the distribution of the other agents' types. Our analysis follows in the same spirit, although we require a slightly weaker notion of dependence. We capitalize on the decision theory literature 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. 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.[8]

This paper considers two distinct cases. In the first, the center receives a signal of its own whose distribution depends on agents' types. In the second, the center does not receive a signal, but agents' types are themselves stochastically dependent. In either case, our construction requires only a very weak form of dependence: for each agent, different values of the agent's type imply different distributions of the center's signal (or of the other agents' types if the center receives no signal of its own). We call this condition stochastic relevance. When stochastic relevance is

[7] See Neeman (2002) for further discussion of mechanism design with correlated information.

[8] See Cooke (1991) and the references therein for a discussion of proper scoring rules.

satisfied, a scoring rule that pays the agent based on the logarithm of the likelihood of the center's information (or the other agents' types if the center receives no signal of its own) is strictly proper, and therefore a suitable scaling of log-likelihood payments induces agents to truthfully reveal their private information.[9] This leads directly to the main results of this paper.

When the center receives a signal of its own whose distribution differs for different types, we show that for any decision rule there exists a balanced transfer scheme that ensures that truthful revelation dominates any announcements that are not close to the truth. When the center does not receive a signal of its own but the distribution of agents' types satisfies a version of stochastic relevance – each agent's type affects the conditional distribution of the other agents' types – dominance or near-dominance results are not possible. Since the center must use the agents' announcements to police them, truthful revelation is not a best response when other agents lie. However, in this case we show that, for any decision rule there is a balanced transfer scheme that implements it in a Bayesian ε-equilibrium.[10,11]

The paper proceeds as follows. Section 2 presents the model. Section 3 derives the results for the case where the center receives a signal of its own. Section 4 presents the result for the case in which the center receives no signal of its own, but the distribution of agents' types satisfies stochastic relevance. Section 5 discusses several extensions, one of which is more fully developed in Appendix A, and Section 6 concludes.

2 The Model and Implementation Strategy

Suppose $N > 1$ agents, indexed by $i = 1, ..., N$, interact with the center. Let G be the set of social alternatives. Although it may be necessary to impose structure on G for certain purposes, such as ensuring the existence of an efficient decision rule, our general results do not require any restrictions on G.

Each agent i has private information or type t_i, where T_i is the set of all possible types for agent i. Following the standard notation we use $t = (t_1, ..., t_N)$ for the vector of types, t_{-i} for the vector

[9] Although we adopt the logarithmic rule in this paper, any strictly proper scoring rule could have been used in our construction. See Cooke (1991) for a discussion of alternative strictly proper scoring rules.

[10] Technically, for budget balance it must be that for every agent i there is another agent i^* such that the conditional distribution of the types of all agents but i and i^* depends on agent i's type.

[11] In a Bayesian ε-equilibrium, the payoff attained by each player's strategy is within ε of the payoff earned by a best response. See d'Aspremont and Gerard-Varet (1982).

of all but player i's type, and t_{-ij} for all but the types of players i and j. The center receives a private signal $z \in Z$. We assume that Z and T_i are compact, convex subsets of finite-dimensional Euclidean spaces with non-empty interiors, and that (z,t) is distributed according to commonly known joint distribution $F(z,t)$.[12] We use $\times T_i$ to denote the product space of the N agents' type spaces.

Each agent's utility is quasilinear in its direct return from the social alternative, g, and money, x, taking the form: $u_i(z,t,g,x) = v_i^*(z,t,g) + x$. The **direct return function** v_i^* depends on the center's information, z, all agents' types, t, and the chosen social alternative, g. We assume that $v_i^*(z,t,g)$ is continuous in t.

A decision rule $g : Z \times (\times T_i) \to G$ maps a type for each agent and the center's information to a social alternative. For simplicity, we assume that $g(z,t)$ is single valued. For $g(z,t)$ that are not single-valued, our implementation result applies to any selection from $g(z,t)$, and therefore this restriction is without loss of generality. A decision rule $g(z,t)$ is (ex post) efficient if and only if the following holds for all states (z,t) and all $\hat{g} \in G$:

$$\sum_i v_i^*(z,t,g(z,t)) \geq \sum_i v_i^*(z,t,\hat{g}). \tag{1}$$

That is, efficient decision rules maximize the aggregate direct return. Given an efficient decision rule, $g(z,t)$, let $V(z,t) = \sum_i v_i^*(z,t,g(z,t))$ be the maximized aggregate direct return function.

Applying the revelation principle, we consider direct revelation mechanisms in which each agent sends a message 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 respectfully 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 $x_i(z,a)$ for each i and a decision rule $g(z,a)$, with the standard interpretation that if the center's information is z and agents announce a, social alternative $g(z,a)$ is realized and transfers $x_i(z,a)$ are made to each i. Although the center's information is not known to the agents when they make their

[12] The dimensionality of the type-spaces may differ for each agent and for the center. Only the existence results in Corollaries 2 and 3 depend on convexity. Otherwise, this assumption can be replaced with R_i being simply connected. Compactness can be relaxed in many places without affecting the results, but at the expense of additional complexity in the mechanism. The assumption of non-emptiness is without loss of generality (except for the case where agents' type spaces are one dimensional), since agents' type spaces can always be defined in terms of the minimal dimension in which the set of feasible types has a non-empty interior.

announcements, the center can verifiably reveal its information at the time transfers are made and the social alternative is implemented.

Denote the vector of transfer rules to all agents by $x(z,a)$, which we call a **transfer scheme**. A transfer scheme is **balanced** if transfers sum to zero: $\sum_{i=1}^{N} x_i(z,a) = 0$. If a decision rule is implemented by a balanced transfer scheme, it does not require an outside subsidy. For fully efficient social choices we are concerned both with maximizing direct returns and with doing so using a balanced transfer scheme.

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 write $v_i^*(z,t,g(z,a))$ as $v_i(z,t,a)$. If there exists a transfer scheme that satisfies a particular solution concept with payoffs $v_i(z,t,a)$, then those transfers implement $g(z,a)$ under that solution concept.

An **announcement strategy** for player i is a function $a_i(t_i) : T_i \to T_i$ that specifies agent i's announcement in the message game as a function of his information. Announcement strategy $a_i(t_i)$ is ε**-truthful** if $||a_i(t_i) - t_i|| < \varepsilon$ for all t_i, where $||\cdot||$ refers to Euclidean distance. When the agent's type is t_i, a single announcement is ε-truthful when $||a_i - t_i|| < \varepsilon$ and ε**-deceptive** if $||a_i - t_i|| \geq \varepsilon$.

Given transfer scheme $x(z,a)$ we say that ε**-deceptions are dominated for player** i, if, for every t_i and a_i' such that $||a_i' - t_i|| \geq \varepsilon$ there exists an announcement $a_i^*(a_i', t_i)$ such that $||a_i^*(a_i', t_i) - t_i|| < \varepsilon$ and for every $a_{-i}(t_{-i})$:

$$\begin{aligned} & E\left\{v_i\left(z,t,a_{-i}(t_{-i}), a_i^*\left(a_i', t_i\right)\right) + x_i\left(z, a_{-i}(t_{-i}), a_i^*\left(a_i', t_i\right)\right) | t_i\right\} \\ > \;& E\left\{v_i\left(z,t,a_{-i}(t_{-i}), a_i'\right) + x_i\left(z, a_{-i}(t_{-i}), a_i'\right) | t_i\right\}, \end{aligned} \tag{2}$$

where the expectation is taken over t_{-i} and z conditional on t_i. That is, ε-deceptions are dominated for player i if for every t_i, any announcement that is ε-deceptive is strictly dominated by one that is ε-truthful. In this case, the set of undominated strategies is a subset of the set of ε-truthful strategies, and we say that **transfer scheme** $x(z,a)$ ε**-implements** $g(z,a)$ **in undominated strategies.**

3 Implementation with Public Information

We begin our analysis by assuming that the center's information serves as a noisy signal of the agents' information, which departs from the Jehiel and Moldovanu (2001) model. By conditioning monetary transfers to the agents on this signal, the center can create incentives for truthful revelation. 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.

Let $f(z|t_i)$ be the density of the center's information conditional on agent i's private information, t_i. Our assumptions on $f(z|t_i)$, assumed to hold for each agent i, are:

Assumption 1: *The conditional densities $f(z|t_i)$ are jointly continuous in z and t_i.*

Assumption 2 (Stochastic Relevance): *The conditional densities $f(z|t_i)$ differ for different values of t_i in the sense that for any distinct types t_i and t_i' there exists a $\hat{z}$ in the interior of Z such that $f(\hat{z}|t_i) \neq f(\hat{z}|t_i')$.*

Taken together, continuity and stochastic relevance imply that $f(z|t_i)$ and $f(z|t_i')$ differ on an open subset of Z and that $f(z|t_i)$ and $f(z|t_i')$ are close together (as functions) if and only if t_i is close to t_i'.

Our final regularity assumption is:

Assumption 3 (Bounded Direct Returns): *There exists $M \geq 0$ such that for all i, t_i and a,*

$$|E\{v_i(z,t,a)|t_i\}| \leq M.$$

The mechanism we propose draws on the decision-theoretic literature on proper scoring functions. The essence of the result is that if it weren't for the direct returns part of an agent's utility, a transfer scheme that paid him according to the log-likelihood of the center's information would induce him to truthfully reveal his private information.

Lemma 1: *Truthful revelation uniquely maximizes the expected log-likelihood function:*[13]

$$t_i = \arg\max_{a_i \in T_i} E_z(\ln f(z|a_i)|t_i).$$

[13] See also Lemma 3.1 in Johnson, Pratt, and Zeckhauser (1990).

Proof of Lemma 1: $E_z(\ln f(z|a_i)|t_i) = \int (\ln f(z|a_i)) f(z|t_i)\,dz$. Since $f(z|a_i)$ is continuous on compact Z, this expectation is bounded above. $E_z(\ln f(z|t_i)|t_i)$ is bounded below since $x \ln x$ is bounded below, and therefore $E_z(\ln f(z|t_i)|t_i)$ is finite. If $E_z(\ln f(z|a_i)|t_i) = -\infty$ for $a_i \neq t_i$, the Lemma is true. For any $a_i \neq t_i$ such that $E_z(\ln f(z|a_i)|t_i)$ is finite,

$$\begin{aligned} E_z(\ln f(z|a_i)|t_i) - E_z(\ln f(z|t_i)|t_i) &= \int (\ln f(z|a_i) - \ln f(z|t_i)) f(z|t_i)\,dz \\ &= \int \left(\ln \frac{f(z|a_i)}{f(z|t_i)}\right) f(z|t_i)\,dz \\ &< \ln \int \frac{f(z|a_i)}{f(z|t_i)} f(z|t_i)\,dz \\ &= \ln \int f(z|a_i)\,dz = 0. \end{aligned}$$

The inequality in the third line follows from Jensen's inequality, strict concavity of the natural logarithm, and stochastic relevance. ■

Since $E_z(\ln f(z|a_i)|t_i)$ is maximized at $a_i = t_i$, a small change in a_i has no first-order effect on the expected transfer. However, since the direct returns function, $v_i(a, z, t)$, is not necessarily maximized at t_i, the agent may have an incentive to deviate from truth-telling to enjoy a personally superior social alternative. Because the center is only able to manipulate agents' incentives through its choice of the transfer function, it will generally not be able to induce the agents to truthfully reveal their multidimensional types.[14] However, since the agents' direct returns from deception are bounded by Assumption 3, the center can pay each agent a large multiple of a log-likelihood payment, swamping the gains from any report that is not arbitrarily close to the truth. The following theorem shows that such a scheme ε-implements any decision rule in undominated strategies and can do so using a balanced transfer scheme.

Theorem 1: *Under Assumptions 1 - 3, for any decision rule there exists a balanced transfer scheme that ε-implements that decision rule in undominated strategies.*

Proof of Theorem: We begin by constructing transfer functions that make ε-deceptions dominated. Later, we ensure the transfer functions are balanced.

[14] Jehiel and Moldovanu (2001) show in the case of independent, mutually-payoff relevant private information that it is only possible to elicit agents' types if a certain integrability condition is satisfied, and that when agents' types are truly multidimensional this condition is satisfied only non-generically.

Let $T_i' = \{(t_i, a_i) \in T_i \times T_i : ||t_i - a_i|| \geq \varepsilon\}$. Note that T_i' is compact and consists of all pairs (t_i, a_i) that cannot be ε-truthful.

For $K > 0$, let

$$x_{iK}(z, a_i) = \max\{-K, \ln f(z|a_i)\}. \tag{3}$$

Function $x_{iK}(z, a_i)$ is a lower truncation of a log-likelihood proper scoring function.

Step 1: For $\delta > 0$ sufficiently small, there exists K sufficiently large that for any i and $(t_i, a_i) \in T_i'$,

$$E\{x_{iK}(z, a_i)|t_i\} \leq E\{x_{iK}(z, t_i)|t_i\} - \delta. \tag{4}$$

Proof of Step 1: Let $L_{iK_i}(t_i, a_i) = E\{x_{iK_i}(z, a_i)|t_i\}$ and $L_i(t_i, a_i) = E\{\ln f(z|a_i)|t_i\}$. By Lemma 1,

$$L_i(t_i, t_i) > L_i(t_i, a_i).$$

Let $\delta_i < \max_{(t_i, a_i) \in T_i'} (L_i(t_i, t_i) - L_i(t_i, a_i))$ and $\delta \leq \min_i \delta_i$. Since $L_{iK_i}(t_i, a_i)$ is a lower truncation of $L_i(t_i, a_i)$,

$$\lim_{K_i \to \infty} L_{iK_i}(t_i, a_i) = L_i(t_i, a_i).$$

Therefore, by compactness of T_i', there exists a uniform K_i such that (4) holds for all $(a_i, t_i) \in T_i'$. Let $K \geq \max_i K_i$. *Q.E.D.*, Step 1.

Choose δ so small and K so large that (4) holds for all i. Let

$$x_i^*(z, a_i) = (2M + 1)\frac{[x_{iK}(z, a_i) + K]}{\delta}. \tag{5}$$

Step 2: If agent i is paid according to $x_i^*(z, a_i)$, ε-deceptions are dominated.

Proof of Step 2: Consider the expected utility reaped by a truthful announcement as compared to announcing a_i with $(t_i, a_i) \in T_i'$.

$$\begin{aligned}
& E\{v_i(z, t, a_{-i}, a_i) + x_i^*(z, a_i)|t_i\} - E\{v_i(z, t, a_{-i}, t_i) + x_i^*(z, t_i)|t_i\} \\
= \ & E\{v_i(z, t, a_{-i}, a_i) - v_i(z, t, a_{-i}, t_i)|t_i\} + E\{x_i^*(z, a_i) - x_i^*(z, t_i)|t_i\} \\
< \ & 2M + (2M+1)\left[\frac{[x_{iK}(z, a_i) + K]}{\delta} - \frac{[x_{iK}(z, t_i) + K]}{\delta}\right] \\
< \ & 2M + (2M+1)\left[-\frac{\delta}{\delta}\right] < 0.
\end{aligned}$$

Payment scheme $x_i^*(z, a_i)$ makes ε-deceptions dominated. Notice that $x_i^*(z, a_i)$ does not depend on a_{-i}. This leaves us significant leeway to balance the budget using the dependence of the full transfer scheme on other players' announcements. If, given a_{-ij}, the transfer to player i, $x_i(z, a_{-ij}, a_j, a_i)$, is independent of a_j, then player j can pay the transfer to player i without affecting j's incentives. *Q.E.D.*, Step 2.

To balance the budget, let $p : N \to N$ be a permutation of the agents with no fixed points with the interpretation that agent i pays transfers to agent $p(i)$. Let

$$x_i(z, a) = x_i^*(z, a_i) - x_{p(i)}^*\left(z, a_{p(i)}\right). \tag{6}$$

Transfer scheme $x_i(z, a)$ makes ε-deceptions dominated and balances the budget. ■

Log-likelihood payment schemes potentially impose arbitrarily large negative penalties on the agents, an infeasible mechanism given agents' finite resources. Step 1 of the proof ensures that there is some truncation of the log-likelihood payment scheme that distinguishes between t_i and a_i when a_i is not ε-truthful. By adding a constant to the truncated scheme, we can provide the same incentives using rewards instead of penalties, as we do in this construction. Of course, since our transfer scheme is balanced, these rewards must be funded by the other agents, and so if agents have limited resources available, this may impair the center's ability to ε-implement certain decision rules. We return to this issue in Section 5.2.

Theorem 1 shows that, for any decision rule, a transfer scheme can be constructed that induces announcements that are arbitrarily close to truthful, provided the center's information is stochastically relevant.[15] The argument we have presented does not depend on continuity of the agents' payoff functions in their announcements. Consequently, it applies even when the decision rule is not continuous and small changes in agents' announcements lead to large changes in the social decision. Discontinuities such as these arise whenever the center's decision involves whether or not to provide a certain good or service, or which of several alternative programs to adopt. Prototypical examples involve whether or not to build a bridge, or which of three locations should be chosen for a new library branch.

Consider the bridge problem. If the government wishes to maximize aggregate welfare, it will build the bridge if and only if the increase in aggregate direct returns from building the bridge are

[15] Matsushima (1990) and Aoyagi (1998) also consider the case of public information, and so Theorem 1 and the other results of this section can be seen as extending Aoyagi (1998) and Matsushima (1990) to the continuous case.

greater than its construction cost. The log-likelihood payment scheme induces agents to report values arbitrarily close to, but not necessarily equal to, their true values. It is therefore possible that an agent who does not want the bridge may underreport his value, and that this deception will result in the bridge not being built even though it is socially desirable to do so. Consequently, individual agents may experience large changes in their utilities. However, for the small deception to influence the center's decision in this way, it must be that the aggregate welfare with and without the bridge is nearly equal. Thus while particular individuals may experience large changes in utility, the change in aggregate welfare will be small.

As we show in Corollary 1 below, this intuition extends beyond the binary outcome case. Whenever the center implements an efficient decision rule, small changes in announcements lead to small changes in aggregate welfare, even though some individuals may experience large changes in direct returns. The reason is that if $g(z,t)$ is an efficient decision rule, maximized aggregate welfare $V(z,t)$ is continuous by Berge's maximum theorem (see Takayama, 1985, p 254) even if $g(z,t)$, itself, is not continuous. This feature, combined with continuity of $v_i^*(z,t,g)$ in t, implies that the loss in aggregate welfare due to small deceptions is small.

Corollary 1: *For any $\delta > 0$, there exist balanced transfers such that for every (z,t), every vector of undominated strategies $a^*(t)$ satisfies:*

$$\left| V(z,t) - \sum v_i^*(z,t,g(z,a^*(t))) \right| < \delta.$$

Proof of Corollary 1: Use the transfers specified in the proof of Theorem 1. For fixed (z,t), let a^* be a generic vector of undominated announcements. Note that

$$\begin{aligned} &\left| V(z,t) - \sum v_i^*(z,t,g(z,a^*)) \right| \\ < \quad &|V(z,t) - V(z,a^*)| + \left| V(z,a^*) - \sum v_i^*(z,t,g(z,a^*)) \right|. \end{aligned} \tag{7}$$

Since $V(z,t)$ is continuous and Z and T_i are compact, $V(z,t)$ is uniformly continuous.[16] Similarly, $\sum v_i^*(z,t,g(z,a^*))$ is uniformly continuous in t and $V(z,a^*) = \sum v_i(z,a^*,g(z,a^*))$. Therefore, there exists $\hat{\varepsilon} > 0$ such that $||a^* - t|| < \hat{\varepsilon}$ implies each of the differences in (7) is made smaller than $\frac{\delta}{2}$. Constructing transfers as in (5) and (6) with $\varepsilon = \frac{\hat{\varepsilon}}{N}$ ensures that all undominated strategies satisfy $||a_i^* - t_i|| < \frac{\hat{\varepsilon}}{N}$ and hence that $||a^* - t|| < \hat{\varepsilon}$, which completes the proof.■

[16] See, for example, Marsden and Hoffman (1993).

3.1 Continuous Direct Returns

Theorem 1 holds quite generally. In particular, Theorem 1 does not require continuity of $v_i^*(z,t,g)$ in any of its arguments.[17] However, while at this level of generality we are able to characterize the set of undominated strategies, without further structure on agents' payoff functions we are unable to characterize the game's Bayesian-Nash equilibria. In this section, we show that if we impose the additional assumption that agents' payoffs are continuous in their *announcements*, then balanced log-likelihood transfers that ensure that all Bayesian-Nash equilibria of the game are ε-truthful can be constructed.[18] Theorem 2 shows that balanced log-likelihood transfers can be constructed under which truthful revelation is an ε-dominant strategy. Corollaries 2 and 3 then establish that balanced log-likelihood payments can be constructed under which an ε-truthful Bayes-Nash equilibrium of the game exists, all Bayes-Nash equilibria of the game are ε-truthful, and realized social outcomes are arbitrarily close to those the center desires.

Since the log-likelihood transfers are continuous, the impact of requiring that payoffs be continuous in announcements is to demand that, given a decision rule, $v_i(z,t,a) = v_i^*(z,t,g(z,a))$ be continuous in a.[19] Continuity of $v_i(z,t,a)$ ensures that the benefits to small deceptions are uniformly small, since continuity of $v_i()$ in a on $\times T_i$ (which is a compact set) implies uniform continuity. Continuity of $v_i()$ in a generally requires that $v_i^*(t,z,g)$ be continuous in g and that decision rule $g(z,a)$ be continuous in a. The latter assumption would be reasonable if, for example, the center is considering what kind of bridge to build (instead of whether to build one at all, as we considered above).[20]

An announcement strategy, $a_i(t_i)$, is **ε-dominant** if, given $\varepsilon > 0$, for every a_i' and any $a_{-i}(t_{-i})$,

$$\begin{aligned} & E\left\{v_i\left(z,t,a_{-i}(t_{-i}),a_i'\right) + x_i^*\left(z,t,a_{-i}(t_{-i}),a_i'\right)|t_i\right\} \\ - \quad & E\left\{v_i\left(z,t,a_{-i}(t_{-i}),a_i(t_i)\right) + x_i^*\left(z,t,a_{-i}(t_{-i}),a_i(t_i)\right)|t_i\right\} \leq \varepsilon. \end{aligned}$$

That is, a strategy is ε-dominant if there is no strategy choice by the other players against which some other strategy outperforms it by more than ε.

[17] Although we have assumed that $v_i^*(z,t,g)$ is continuous in t, Theorem 1 does not require this assumption.

[18] This additional structure is useful in the following section, which considers the case where agents types are correlated but the center has no information of its own.

[19] This rules out discrete decisions such as whether or not to build a bridge.

[20] In addition, for many discontinuous decision rules our mechanism can be used to realize a continuous approximation of the discontinuous rule (e.g., a rule that specifies the probability of building the bridge), with approximate realization of the approximate decision rule coming arbitrarily close to realizing the true decision rule.

Theorem 2: *Suppose $f(z|t_i)$ satisfies Assumptions 1 and 2 and direct returns functions $v_i(z,t,a)$ are continuous in a. Then there exists a balanced transfer scheme under which truthful revelation is an ε-dominant strategy for each i.*

Proof of Theorem 2: Continuity of direct returns on compact $\times T_i$ implies Assumption 3. Therefore, for any $\varepsilon^* > 0$, we can construct payments according to (3) and (5) in the proof of Theorem 1 such that truthful revelation dominates any announcement with $||a_i - t_i|| \geq \varepsilon^*$. Without affecting incentives, payments can be balanced according to (6).

Continuity of $v_i(z,t,a)$ and $x_i(z,a)$ in a on compact $\times T_i$ imply that expected utility is uniformly continuous in a. Hence for any $\varepsilon > 0$ there exists ε^* such that for every i, $a_{-i}(t_{-i})$, and t_i, if $||a_i - t_i|| < \varepsilon^*$, then

$$\begin{aligned} & E\left\{v_i\left(z,t,a_{-i}\left(t_{-i}\right),a_i\right)+x_i^*\left(z,a_{-i}\left(t_{-i}\right),a_i\right)|t_i\right\} \\ - \quad & E\left\{v_i\left(z,t,a_{-i}\left(t_{-i}\right),t_i\right)+x_i^*\left(z,a_{-i}\left(t_{-i}\right),t_i\right)|t_i\right\} < \varepsilon \end{aligned}$$

Since by Claim 2 any announcement a_i' with $||a_i' - t_i|| \geq \varepsilon^*$ is dominated by truthful revelation, this completes the proof. ■

Under the construction in Theorem 2, truthful revelation is arbitrarily close to being a dominant strategy. However, truthful revelation need not be a best response to the strategies chosen by the other players. Nevertheless, because payoffs are continuous, strategy sets are compact, and ε-deceptions are dominated, the players' best-response correspondences are well defined, and for any t_i the set of best responses for player i to opponents' strategies $a_{-i}(t_{-i})$ must be ε-truthful. This allows us to characterize the set of Bayes-Nash equilibria under the log-likelihood transfer scheme.

An **announcement vector** $a(t)$ **is a Bayes-Nash equilibrium given transfer scheme** $x(z,a)$ if for each i, t_i, and a_i' :

$$\begin{aligned} & E\left\{v_i\left(z,t,a_{-i}\left(t_{-i}\right),a_i\left(t_i\right)\right)+x_i^*\left(z,a_{-i},a_i\left(t_i\right)\right)|t_i\right\} \\ \geq \quad & E\left\{v_i\left(z,t,a_{-i}\left(t_{-i}\right),a_i'\right)+x_i^*\left(z,a_{-i}\left(t_{-i}\right),a_i'\right)|t_i\right\}. \end{aligned}$$

Corollary 2: *Suppose $f(z|t_i)$ satisfies Assumptions 1 and 2 and direct returns functions $v_i(z,t,a)$ are continuous in a and quasiconcave in a_i. There exists a balanced transfer scheme under which:*

i) *There is a pure strategy Bayes-Nash equilibrium in which agents play ε-truthful strategies.*

ii) *In any pure-strategy Bayes-Nash equilibrium all players play ε-truthful strategies.*

Proof of Corollary 2: Assumption 3 follows from the continuity of $v_i(z,t,a)$ on compact $\times T_i$. Use the transfers specified in (3), (5) and (6). Existence is standard in this game since strategy sets are compact, convex, non-empty subsets of Euclidean space and payoffs are continuous in a and quasiconcave in a_i. See Fudenberg and Tirole, 1993, Theorem 1.2. Part ii) follows from constructing transfers that make announcements that are not ε-truthful strictly dominated for agent i, and noting that in any Nash equilibrium, no strictly dominated strategy is played with positive probability. ■

Corollary 2 establishes that a Bayes-Nash equilibrium exists under the log-likelihood transfer scheme and that the transfer scheme can be constructed to make all Bayes-Nash equilibria ε-truthful. Corollary 3 continues to exploit continuity in order to establish that given a decision rule, balanced log-likelihood transfers can be constructed under which the outcomes realized in any Bayes-Nash equilibrium of the mechanism are arbitrarily close to those dictated by the decision rule.

Corollary 3: *Suppose $f(z|t_i)$ satisfies Assumptions 1 and 2. Let $g(z,a)$ be continuous in a, and suppose direct returns functions $v_i^*(z,t,g(z,a))$ are continuous in a and quasiconcave in a_i. For any $\hat{\varepsilon} > 0$ there exists a balanced transfer scheme under which every Bayes-Nash equilibrium strategy $a^*(t)$* satisfies for every (z,t) :

i) $||g(z,a^*(t)) - g(z,t)|| < \hat{\varepsilon}$, *and*

ii) $|\sum v_i^*(z,t,g(z,a^*(t))) - \sum v_i^*(z,t,g(z,t))| < \hat{\varepsilon}$.

Proof of Corollary 3: Continuity of $v_i(z,t,g)$ and $g(z,t)$ on compact set $\times T_i$ implies uniform continuity and Assumption 3. By uniform continuity, the transfers specified in (3), (5), and (6) can be constructed so that i) holds and $|v_i^*(z,t,g(z,a^*(t))) - v_i^*(z,t,g(z,t))| < \frac{\hat{\varepsilon}}{N}$ for each i, from which part ii) follows. ■

Before continuing, a few comments about Corollaries 1, 2, and 3 are in order. First, Corollary 3 holds for any decision rule, whereas Corollary 1 is a similar result that holds only for efficient decision rules. The difference is that, in the absence of continuity in announcements, Corollary 1 relies on the theorem of the maximum to ensure that changes in aggregate welfare are small, and the theorem of the maximum only applies if the decision rule is the solution to some optimization problem.[21] Corollary 3 assumes continuity in announcements directly, and derives a stronger result. Second, the assumption of quasiconcavity in Corollaries 2 and 3 can be relaxed. If all assumptions but quasiconcavity are satisfied, the results continue to hold, but the equilibria may be in mixed strategies (see Fudenberg and Tirole, 1993, Theorem 1.3 for existence). In this case, $a_i(t_i)$ are interpreted as referring to probability distributions over T_i, and the statements apply to almost all announcements in the support of these mixed strategies.

4 Implementation with Correlated Types

The implementation results we have presented so far employ a multiple of log-likelihood payments to overwhelm the direct returns to distortion. For a large enough multiple, the agent cares almost exclusively about the transfers rather than about direct returns, and this induces him to make a nearly truthful announcement. The same technique can be applied when the center gets no information (or the center's information is independent of the agents'), but the distribution of agents' $j \neq i$ types conditional on agent i's type satisfies Assumption 1 and a version of stochastic relevance adapted to the present context (defined below). In this case, we can show that if the direct returns to player i vary continuously with his announcement, then there exist transfer schemes that make truthful revelation by every player a Bayesian ε-equilibrium.

For this section, we assume that the center gets no private signal and we omit argument z from the direct returns and transfer functions. We investigate instead how the center may use other agents' announcements to police agent i when types are stochastically related.

We begin by defining Bayesian ε-equilibrium. Suppose that players $j \neq i$ use announcement strategies $a_{-i}(t_{-i})$. Given $\varepsilon > 0$, announcement $a_i(t_i)$ is an **ε-best response for player i to $a_{-i}(t_{-i})$ given transfer functions** $x_i(a)$ if for all t_i and $\hat{a}_i \in T_i$:

$$\begin{aligned} & E\left\{v_i\left(t, a_{-i}\left(t_{-i}\right), a_i\left(t_i\right)\right) + x_i^*\left(a_{-i}\left(t_{-i}\right), a_i\left(t_i\right)\right) | t_i\right\} \\ \geq\; & E\left\{v_i\left(t, a_{-i}\left(t_{-i}\right), \hat{a}_i\right) + x_i^*\left(a_{-i}\left(t_{-i}\right), \hat{a}_i\right) | t_i\right\} - \varepsilon. \end{aligned} \tag{8}$$

That is, $a_i(t_i)$ is an ε-best response if no other strategy outperforms it by more than ε. A vector of announcement strategies, $a_i(t_i)$, comprises a **Bayesian ε-equilibrium given transfer scheme**

[21] Indeed, a version of Corollary 1 would hold whenever $g(z, a)$ is derived as the solution to some optimization problem, not just the aggregate welfare maximization problem.

$x(a)$ if (8) is satisfied for each i. Truthful revelation is a Bayesian ε-equilibrium given $x(a)$ if (8) is satisfied for each i when $a_j(t_j) = t_j$ for all j.

In Theorems 1 and 2, agents' announcements did not affect the center's information. In the next result, the center receives no signal of its own, but uses the announcements of players $j \neq i$ in order to induce truthful revelation by player i. To balance the budget we assign a player i^* to pay the truth-inducing transfers to player i. The following strengthening of Assumption 2 ensures that this can be done without altering the incentives provided to player i^*.

Assumption 4: *Suppose $N \geq 3$. For each i,* there exists an agent $q(i) \neq i$ such that $f\left(t_{-iq(i)}|t_i\right)$ *satisfies the stochastic relevance property. That is, for any distinct types t_i and t_i' there exists $t_{-iq(i)}$ such that:*

$$f\left(t_{-iq(i)}|t_i\right) \neq f\left(t_{-iq(i)}|t_i'\right).$$

Agent $q(i)$ is the agent who will balance the budget with respect to agent i. In order to provide incentives to agent i, agent i will be asked to predict the distribution of $t_{-iq(i)}$, the reports of all agents other than i and $q(i)$. The need for at least one such agent leads to the requirement that $N \geq 3$. It need not be the case that $q(i)$ is distinct for each i. Further, Assumption 4 is satisfied if, for each agent i, there is another agent whose type is stochastically relevant for agent i's type.

Theorem 3: *Suppose $f(\cdot|t_i)$ satisfies Assumptions 1 and 4. If $v_i(t,a)$ are continuous in a, then there exist transfers that make truthful revelation by each player a Bayesian ε-equilibrium.*

Proof of Theorem 3: The proof follows the proof of Theorem 2. Assumption 3 follows from continuity of $v_i(t,a)$ in a on compact $\times T_i$. Let $a_i^*(a_{-i}(t_{-i}), t_i)$ denote player i's best response to strategies $a_{-i}(t_{-i})$ given t_i. Suppose players $j \neq i$ announce truthfully: $a_{-i}(t_{-i}) = t_{-i}$. Choose $q(i)$ according to Assumption 4. For each i, let $z_i = a_{-iq(i)}\left(t_{-iq(i)}\right) = t_{-iq(i)}$ play the role of z in (3), (5), and (6). The arguments in Claims 1 and 2 show that for any $\varepsilon^* > 0$, transfers can be constructed so that any strategy $a_i(t_i)$ that satisfies $||a_i(t_i) - t_i|| \geq \varepsilon^*$ for some t_i is dominated by one that reveals truthfully, conditional on $t_i = t_i$.

As in the proof of Theorem 2, since direct returns and transfers are continuous in a on compact $\times T_i$, each agent's expected utility is uniformly continuous in a. Therefore, for any $\varepsilon > 0$ there exists $\varepsilon^* > 0$ such that if $||a_i - t_i|| < \varepsilon^*$ then

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

Hence truthful revelation is an ε-best response, provided that all other players announce truthfully. Since i was chosen arbitrarily, truthful revelation is a Bayesian ε-equilibrium. Budget balance follows from the same construction used in the proof of Theorem 1: agent j pays transfers to any agent i for which $j=q\left(i\right)$. By construction, these transfers do not depend on j's announcement. ■

Assumption 4 is essentially the same as Assumption 2 in Aoyagi (1998). As Aoyagi notes in the finite case, Assumption 2 is stronger than the d'Aspremont-Gérard-Varet (1979,1982) compatibility condition since compatibility can be satisfied by independent types. However, d'AGV compatibility applies to efficient decision rules when payoffs are not mutually payoff-relevant. Aoyagi's Assumption 2 applies to any decision rule when payoffs may be mutually payoff-relevant, although unlike in the present paper, Aoyagi's results are restricted to the finite-type case.

Aoyagi observes as well that the general relation between Stochastic Relevance, his Assumption 2, and d'AGV Compatibility is likely to be quite complex, and we are more concerned with demonstrating the robustness of our log-likelihood construction than with pinpointing necessary conditions for (approximate) incentive compatibility. However, it is worth noting that Jehiel and Moldovanu (2001) shows that incentive compatibility is generically impossible when types are independent, and Assumption 4 is a very weak form of dependence. Hence if Assumption 4 is not necessary, the necessary condition must fall between independence and stochastic relevance, and it is unclear to us whether the space between these two concepts is economically meaningful.

5 Extensions

As should be clear from the constructions in Theorems 1 and 3, the scoring-rule based approach to implementation is very robust; very little structure is needed to ensure implemetability. In this section, we illustrate how the basic mechanism can be extended to address cases where agents fact ex ante participation constraints, interim participation constraints, or limited liability constraints. We also discuss how our techniques can be extended to the case considered by Johnson, Pratt, and Zeckhauser (1990, henceforth JPZ), in which agents may take costly actions in addition to making announcements. To better match the literature on Bayesian mechanism design, we present the extensions in the environment where the center does not receive a signal but the agents' types satisfy stochastic relevance.

5.1 Ex Ante Participation Constraints

We begin by considering the case where each agent must expect to break even ex ante. Because our mechanism balances for each realization of agents' types, this requirement poses little additional challenge. For fixed decision rule $g(a)$, we assume that ex ante expected surplus is positive, i.e., that, assuming that agents announce truthfully, the project is expected to be worthwhile:

$$E_t\left(\sum_i v_i^*(t, g(t))\right) \geq 0.$$

Let $x_i(t)$ be defined as in the proof of Theorem 3. Under these transfers, truthful reporting is a Bayesian ε-equilibrium. Agent i's ex ante expected utility is:

$$k_i \equiv E_t\left(v_i^*(t, g(t)) + x_i(t)\right).$$

Suppose that, prior to participating in the mechanism, each agent is charged participation fee $k_i - \frac{\sum_{j=1}^N k_j}{N}$. Since this fee does not depend on the agents' announcements, truthful revelation remains an Bayesian ε-equilibrium. Ex ante expected utility for agent i is then:

$$\begin{aligned} & E_t\left(v_i^*(t, g(t)) + x_i(t)\right) - k_i + \frac{\sum_{j=1}^N k_j}{N} \\ = & \frac{\sum_{j=1}^N k_j}{N} = \frac{\sum_{j=1}^N}{N} E_t v_j^*(t, g(t)) \geq 0. \end{aligned}$$

Hence the ex ante participation constraint is satisfied for each agent.

5.2 Interim Participation Constraints, Full Surplus Extraction, and Limited Liability

McAfee and Reny (1992) show 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. To summarize their mechanism, let $\pi_i(t_i)$ be the expected rent earned by agent i when his type is t_i. If $\pi_i(t_i)$ is continuous, then, when agents' types are correlated, it is possible to choose a finite set of participation fee schedules $z_1(t), ..., z_M(t)$ such that

$$0 \leq \pi_i(t_i) - \min_{1 \leq m \leq M} \int z_m(t_{-i}) f(t_{-i}|t_i)\, dt_{-i} < \varepsilon$$

for each $t_i \in T_i$.[22] If the center offers the agent a choice among the M fee schedules before the subsequent game is played, and if for each t_i agent i chooses the schedule $z_m(t)$ that maximizes his expected surplus and that choice is not used in the subsequent game, then by offering the agent this menu of fee schedules the center can extract all of the agent's rents and can also ensure that the agent's expected rent from participating in the mechanism is non-negative, i.e., that interim participation constraints can be satisfied.

While McAfee and Reny (1992) show 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, in general, that such a mechanism exists when agents have multidimensional, continuous types and interdependent valuations.

We show that if agents' types are correlated, then any decision rule can be implemented arbitrarily closely. This, coupled with the McAfee-Reny result, establishes that the center can extract 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.

One of the criticisms of the correlated-mechanism-design literature is that the mechanisms rely on payments that become very large as agents' types become "nearly independent," and therefore become problematic if agents have limited resources. While this remains true in the present case, this problem can also be addressed by coupling our mechanism with the McAfee-Reny mechanism. In particular, suppose that instead of having agents pay other agents to balance the mechanism, we instead had the center make all incentive payments to the agents. In this case, the ability to implement a particular decision rule would not depend on agents' limited resources. Although such payments would result in rents being transferred to the agents, a menu of McAfee-Reny-style participation fees could extract this surplus without violating agents' participation constraints.

[22] See McAfee and Reny (1992), equation 1.2. Since stochastic relevance and Lemma 1 imply that for every t_i there is a function of all agents' types taking a unique minimum at t_i, stochastic relevance is sufficient for full surplus extraction (see Remark 2, p. 406).

5.3 Costly Actions by Agents

Our analysis also extends readily to the environment considered in JPZ, in which, in addition to sending messages, agents may also take actions before the center acts that affect the welfare of all agents. The informational requirements are unchanged. Either the center must receive a stochastically relevant signal of agents' types or else for each agent there is stochastic relevance for the actions of at least one other agent. If either requirement is satisfied, the center may employ log-likelihood payments to extract the agents' private information and then use that information to create ex post incentives that align the agents' uncoordinated actions with the desired decision rule.

The simplest case of costly actions arises when there are costs associated with sending the messages themselves. Formally, this situation could be modeled by letting the agents' direct returns functions depend on the agents' messages directly, as in $v_i(z, t, g, a)$. If direct returns are bounded, the arguments of Theorems 1, 2, and 3 go through essentially unmodified. However, in adopting this approach, applying the revelation principal incurs some loss of generality, since doing so restricts the agent's action space to be identical to his type space.

Our results also extend to two cases in which agents may take costly, uncoordinated actions, distinct from the messages they send. The first corresponds to JPZ's "responsive actions" environment, in which there is a one-to-one relationship between agents' types and desired actions. In this case, actions can be used as messages and our implementation results follow. The second corresponds to the case in which the center desires more than one type of agent to take the same action. In this case, the space of observed actions is not rich enough to identify agents's types. Our implementation results once again follow, but in this case the agents' actions must be augmented with additional messages to the center.[23]

6 Conclusion

This paper extends the mechanism design literature to show the possibility of implementing practically any decision rule when agents' types are continuous, multidimensional, and mutually payoff relevant, provided that either the center receives an informative signal of agents' types or agents' types are correlated. Thus we provide a converse to Jehiel and Moldovanu (2001), who show that if

[23] For the sake of brevity, we relegate the formal statement and proofs of these claims to the Appendix.

agents' types are independent, then efficient design is generically not possible in this environment. Further, we complement the analysis of McAfee and Reny (1992) by showing that there is an ex post efficient mechanism in the multidimensional, continuous, mutually payoff-relevant case.

The scoring-rule based approach we adopt has the advantage of being simpler than those commonly adopted in the mechanism design literature. Stochastic relevance 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 (1979; 1982), the linear independence condition of Crémer and McLean (1985; 1988), or the generalization of the Crémer-McLean condition found in McAfee and Reny (1992), which must hold for all prior distributions for each agent's type. In addition to being simpler, stochastic relevance is also slightly weaker than these conditions. The log-likelihood payments used in our mechanism are also relatively simple to construct and our proof provides a blueprint for doing so. This is in contrast to the approach adopted by other authors, who generally prove the existence of a mechanism but provide little guidance as to how it should be constructed.[24]

Another advantage of our approach is that it is robust. We address the problem of implementing decision rules when agents' types are continuous, multidimensional, and mutually payoff-relevant because this problem has not been previously solved and presents the greatest technical challenge. However, our methods would also apply to the case of finite types. In addition our methods extend well beyond the traditional mechanism design problem. For example, as illustrated in the Appendix, our techniques can also be applied to the case when, in addition to having mutually payoff-relevant private information, agents may also take actions which themselves may create real externalities.

References

[1] Aoyagi, M. (1998): "Correlated Types and Bayesian Incentive Compatible Mechanisms with Budget Balancedness," *Journal of Economic Theory*, 79, 142-151.

[2] Clarke, E. (1971): "Multipart Pricing of Public Goods," *Public Choice*, 8, 19-33.

[24] Freuqently, such approches rely on a linear systems approach to demonstrate existence. See d'Aspremont, Crémer, and Gérard-Varet (1990) for a survey of the use of this method.

[3] Cooke, Roger M. (1991): *Experts in Uncertainty: Opinion and Subjective Probability in Science.* New York: Oxford University Press.

[4] Crémer, J., and R. McLean (1985): "Optimal Selling Strategies Under Uncertainty for a Discriminating Monopolist When Demands Are Interdependent," *Econometrica*, 53, 345-361.

[5] ________________________(1988): "Full Extraction of Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica,* 56, 1247-1257.

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

[7] ________________________(1982): "Bayesian Incentive Compatible Beliefs," *Journal of Mathematical Economics*, 10, 83-103.

[8] ________________________(1998): "Linear Inequality Methods to Enforce Partnerships Under Uncertainty: An Overview," *Games and Economic Behavior*, 25, 311-336.

[9] 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.

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

[11] Fudenberg, D. and J. Tirole (1993): *Game Theory,* MIT Press, Cambridge, MA.

[12] Green, J. and J.-J. Laffont (1977): "Characterization of Satisfactory Mechanisms for the Revelation of Preferences for Public Goods," Econometrica, 45, 427-438.

[13] Green, J. and J.-J. Laffont (1979): *Incentives in Public Decision Making.* Amsterdam: North Holland.

[14] Groves, T. (1973): "Incentives in Teams," *Econometrica*, 41, 617-663.

[15] Jackson, M. (2001): "A Crash Course in Implementation Theory," Discussion Paper, Caltech, forthcoming in *Social Choice and Welfare.*

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

[17] Johnson, S., J. Pratt, and R. Zeckhauser (1990): "Efficiency Despite Mutually Payoff-Relevant Private Information: The Finite Case," *Econometrica,* 58, 873-900.

[18] Krishna, V. (2002): *Auction Theory.* London: Academic Press.

[19] Marsden, J. and M. Hoffman (1993): *Elementary Classical Analysis,* Second Edition. New York: W.H. Freeman and Co.

[20] Mas-Colell, A., M. Whinston, and J. Green (1995): *Microeconomic Theory.* New York: Oxford University Press.

[21] Maskin, E. (1992): "Auctions and Privatization," in H. Siebert (ed.), *Privatization,* Kiel: Institut fur Weltwirtschaften der Universität Kiel, 115-136.

[22] ____________________(1990): "Dominant Strategy Mechanisms with Mutually Payoff-Relevant Information and with Public Information," *Economics Letters,* 34, 109-112.

[23] ____________________(1991): "Incentive Compatible Mechanisms with Full Transferability," *Journal of Economic Theory,* 54, 198-203.

[24] McAfee, P. and P. Reny (1992): "Correlated Information and Mechanism Design," *Econometrica,* 60, 395-421.

[25] McLean, R. and A. Postlewaite (2003): "Informational Size and Efficient Auctions," *Review of Economic Studies,* forthcoming.

[26] Neeman, Z. (1999): "The Relevance of Private Information in Mechanism Design," Discussion Paper, Boston University (June 28, 1999).

[27] Takayama, A. (1985): *Mathematical Economics*, Second Edition. Cambridge, UK: Cambridge University Press.

[28] Vickrey, W. (1961): "Counterspeculation, Auctions, and Competitive Sealed-Tenders," *Journal of Finance*, 16, 8-37.

A Costly Actions by Agents

Thusfar we have framed this paper in the style of the mechanism design literature and presented the results as complementing the Jehiel and Moldovanu (2001) impossibility results. However, our analysis extends readily to the substantially more general environment considered in JPZ. Hence, the present work can also be seen as extending JPZ to the continuous case.

In the standard mechanism design problem, agents' merely send messages that have no direct effect on their payoffs. It is only through the center's decision rule that an agent's announcement affects his payoff. While in many circumstances, such as auctions, this is natural, there are many collective decision problems in which, before the center acts, an agent takes actions that affect his own welfare and possibly that of others.

To illustrate, consider a group of firms who are polluting some resource. Their actions affect each other, i.e., convey the usual negative externality. Beyond this, the members will be subject to government regulation that is crafted in response to information gleaned from the polluters. Each polluter knows more than his peers or the regulator about the composition of his discharges, and hence the magnitude of their consequences for others. In such cases, each agent will have an incentive to take actions and send messages that influence the social alternative chosen, e.g., to dump excessively to make it appear that it will be expensive for him to cut back, securing him greater leeway under regulation. Such are the nature of the problems considered by JPZ, although they deal with finitely many types.

Our log-likelihood mechanism can perform effectively in cases where, beyond eliciting honest revelation of types, agents' costly actions must also be appropriately aligned. The informational requirements are unchanged. Either the center must receive a stochastically relevant signal of agents' types or else for each agent there is stochastic relevance from the actions of at least one other agent. If either requirement is satisfied, the center may employ log-likelihood payments to extract the agents' private information and then use that information to create ex post incentives that align the agents' uncoordinated actions agree with the desired decision rule.[25] Thus, in the preceding example, the regulatory authority would anticipate the exaggerated dumping, and would make the incentives sufficiently strong that each would dump as he would in a fully coordinated

[25] The actions are uncoordinated in the sense that agents choose them before the center acts. Each agent bases his action choice on his private information only, and while the center may provide incentives for such actions, it cannot directly control or coordinate them.

game.

The simplest case of costly actions arises when there are costs associated with sending the messages themselves. Formally, this situation could be modeled by letting the agents' direct returns functions depend on the agents' messages directly, as in $v_i(z, t, g, a)$. If direct returns are bounded, the arguments of Theorems 1, 2, and 3 go through essentially unmodified. However, in adopting this approach, applying the revelation principal incurs some loss of generality, since doing so restricts the agent's action space to be identical to its type space. While this is sometimes a reasonable restriction (see Theorem 4 below), often it is not. Nevertheless, the intuition behind our main results is still relevant. When direct returns are bounded, log-likelihood payments can induce agents to make nearly truthful announcements, even when the act of announcing is costly.

To further explore the connection between our present work in the continuous case and JPZ, we now consider two cases in which agents may take costly, uncoordinated actions, distinct from the messages they send. The first corresponds to JPZ's "responsive actions" environment, in which there is a one-to-one relationship between agents' types and desired actions. In this case, actions can be used as messages and our implementation results follow. The second corresponds to the case in which the center desires more than one type of agent to take the same action. In this case, the space of observed actions is not rich enough to identify agents's types. Our implementation results once again follow, but in this case the agents' actions must be augmented with additional messages to the center.

Since the construction in Theorem 1 is the driving force behind all results, for the sake of brevity we present in this section analogues to Theorem 1 for each of the two cases under consideration. Extensions of the other results to these cases follow the constructions in the earlier part of the paper.

We begin by defining additional notation. Let B_i be the set of actions available to agent i, and let b_i be a typical element of that set. Thus, for example, B_i may be all the possible ways a polluter can dump its waste prior to being regulated. Importantly, we do not require that B_i be of the same dimension as agent i's type space, T_i. Actions may be more or less complex than types.

Let $b = (b_1, ..., b_N)$ refer to the vector of actions taken by all agents. In order to account for the mutually payoff-relevant direct returns from these actions, we integrate them into agents' direct returns functions and write $v_i^* (z, t, g, b)$. Since the center observes b before making its decision, a decision rule may now also depend on b. Let such a decision rule be written as $g(z, t, b)$.

Assumption $3'$ extends the bounded returns assumption to this case.

Assumption $3'$: *There exists* $M \geq 0$ *such that for all* i*,* t_i*,* g*,* and b*,*

$$\left| E\left\{ v_i^* \left(z, t, g, b\right) | t_i \right\} \right| \leq M.$$

In addition to its decision rule, the center must now also consider the agents' uncoordinated actions. Each agent chooses b_i knowing only t_i. Let $b_i\left(t_i\right)$ denote this dependence. Let $b\left(t\right) = \left(b_1\left(t_1\right), ..., b_N\left(t_N\right)\right)$ denote action rules for each agent.

We begin with the simplest case. Following JPZ, we say an action rule is responsive if for each i, $b_i\left(t_i\right) = b_i\left(t_i'\right) \Rightarrow t_i = t_i'$. That is, an action rule is responsive if different types choose different actions. If responsiveness holds, then an agent's action choice is effectively an announcement of a type. If direct returns are bounded, log-likelihood payments can be used to induce the agent to choose an action for which the associated type is arbitrarily close to his true type.

Define the inverse action function $\rho_i\left(b_i\right)$ as follows:

$$\rho_i\left(b_i\right) = \begin{cases} b_i^{-1}\left(b_i\right) & \text{if } b_i\left(t_i\right) = b_i \text{ for some } t_i \in T_i \\ t^* & \text{otherwise} \end{cases},$$

where t^* is chosen so that $||t^* - t_i|| \geq T$, where $T = \max ||t_i - t_i'||$ for all i,and t_i and $t_i' \in T_i$. Such an t^* exists by compactness of T_i. The idea is that taking an action that should not be taken by any type of agent given action rule $b_i\left(t_i\right)$ is interpreted as "announcing" a type that is far from every $t_i \in T_i$.

Suppose the center employs the following mechanism. Before play, it announces decision rule $g\left(z, b\right)$, desired action rule $b\left(t\right)$ and transfer scheme $x\left(z, b\right)$. After observing their types, agents choose actions b_i. The center then implements decision $g\left(z, b\right)$ and makes transfers $x\left(z, b\right)$. Following the definition of ε-truthful undominated strategies in section 2, we say that action b_i is an ε-truthful action given $b_i\left(t_i\right)$ if $||\rho_i\left(b_i\right) - t_i|| < \varepsilon$ when the agent's true type is t_i. An action is ε-deceptive given $b_i\left(t_i\right)$ if $||\rho_i\left(b_i\right) - t_i|| \geq \varepsilon$. We say that ε-deceptive actions are dominated for player i given $b_i\left(t_i\right)$ if for every t_i and b_i' such that $||\rho_i\left(b_i'\right) - t_i|| \geq \varepsilon$ there exists an action $\hat{b}_i\left(b_i', t_i\right)$ such that $\left|\left|\rho_i\left(\hat{b}_i\left(b_i', t_i\right)\right) - t_i\right|\right| < \varepsilon$ and for all t_{-i} and b_{-i}:

$$\begin{aligned} & E\left\{ v_i\left(z, t, g\left(z, t, \left(\hat{b}_i\left(b_i', t_i\right), b_{-i}\right)\right), \left(\hat{b}_i\left(b_i', t_i\right), b_{-i}\right)\right) + x_i\left(z, \left(\hat{b}_i\left(b_i', t_i\right), b_{-i}\right)\right) | t_i \right\} \\ > \ & E\left\{ v_i\left(z, t, g\left(z, t, \left(b_i', b_{-i}\right)\right), \left(b_i', b_{-i}\right)\right) + x_i\left(z, \left(b_i', b_{-i}\right)\right) | t_i \right\}. \end{aligned}$$

That is, ε-deceptive actions are dominated if, given the action rule, for any ε-deceptive action there is an ε-truthful action that dominates it. When ε-deceptive actions are dominated for each player, we say that the transfer scheme ε-implements the decision *and* action rules in undominated strategies.

Theorem 4: *Under Assumptions 1, 2, and 3′, for any decision rule and responsive action rule there exists a balanced transfer scheme that ε-implements them in undominated strategies.*

Proof of Theorem 4: Given responsiveness, the decision rule $g(z, b)$ can be written $\hat{g}(z, a)$, where $a_i = \rho_i(b_i)$. Treat action b_i as if the agent announced $a_i = \rho_i(b_i)$ and apply the construction in Theorem 1. Transfers satisfying (3), (5), and (6) are balanced and ε-implement $\hat{g}(z, a)$ in undominated strategies, which is equivalent to ε-implementing $g(z, b)$ and $b(t)$ in undominated strategies.

The responsive case is straightforward. Since agents types are encoded in their actions, actions effectively announce types. Thus we are able to apply the revelation principle after translating agents actions into the corresponding types according to the announced action rule. As in our earlier analysis, continuity is a concern. If $b(t)$ is not continuous, then nearby types may be asked to choose very different actions, and small deceptions may lead to large changes in individual payoffs. Nevertheless, if the center chooses $b(t)$ and $g(z, a, b)$ to maximize expected aggregate direct returns, then the analogue to Corollary 1 shows that changes in expected aggregate direct return will be small, even if some individuals experience large changes in their direct returns.[26]

When the center's desired decision rule is not responsive, then actions no longer act to announce types. Nevertheless, implementation remains possible using additional messages. As long as the center receives a stochastically relevant signal, it can induce agents to announce their types with arbitrary precision. Once the center has gleaned this information, it is a simple matter to check whether the action the agent has taken agrees with the type he has announced. Imposing a sufficiently large penalty if action and announcement fail to agree ensures that the agent will announce (nearly) truthfully and choose the action appropriate for his announcement.

[26] Since agents' actions are chosen after learning r_i but before the center acts, ex post efficiency is no longer the appropriate efficiency concept. Rather, the center should choose $g(z, r, b)$ and $b(r)$ to maximize ex ante expected direct returns.

Since actions no longer provide sufficient information, we consider mechanisms where the center announces a decision rule $g(z,a,b)$, action rule $b(t)$, and vector of transfer rules $h_i(z,a,b)$ that depend on the center's information and agents' messages and actions. Agents send messages $a_i \in T_i$ and choose actions $b_i \in B_i$, following which the center implements the decision and makes transfers.

Given action rule $b_i()$, we say that agent i's action **agrees** with his announcement if $b_i = b_i(a_i)$. Since the center observes the agent's announcement and action, it is straightforward to induce the agent to choose them such that they agree. By Assumption $3'$, the maximum gain from choosing a non-agreeing action is bounded by $2M$. Therefore, penalizing the agent $2M$ if his action and announcement do not agree ensures that any action-announcement pair that does not agree is dominated by one that does.

Theorem 5: *Under Assumptions 1, 2, and $3'$, for any decision rule and action rule there exists a balanced transfer scheme that ε-implements the decision rule in undominated strategies and ensures agents' actions agree with their announcements.*

Proof of Theorem 5: Transfers constructed according to (3), (5), and (6) are balanced and ε-implement the decision rule in undominated strategies. To induce agreement, augment $x_i(z,a)$ as follows. Let

$$s_i^*(a_i,b_i) = \begin{cases} 0 & \text{if } b_i = b_i(a_i) \\ -2M & \text{otherwise} \end{cases}.$$

Let $s_i(a,b) = s_i^*(a_i,b_i) - \sum_{j \neq i} \frac{s_j(a_j,b_j)}{N-1}$. Given a_i, balanced transfers $s_i(a,b)$ ensure that agent i's announcement and action agree. Letting i's total transfer be $h_i(z,a,b) = x_i(z,a) + s_i(a,b)$ satisfies both requirements of the theorem.[27]

Theorems 4 and 5 extend Theorem 1 to the continuous version of the problem considered in JPZ and show that mutually-payoff relevant, multi-dimensional private information may be successfully elicited and appropriate actions induced. Similar extensions can be developed for the remaining results in Section 3. As in all results in this paper, the key lies in using log-likelihood transfers to induce nearly truthful revelation. To the extent that the desired actions are responsive, the center's mechanism can use the actions themselves in lieu of messages, as in Theorem 4. If actions

[27] Note that in equilibrium agents choose actions that agree with their announcements and no penalties are assessed.

are not responsive, then the center can be successful by adopting a different strategy, using log-likelihood payments to induce nearly truthful revelation and then penalizing the agent if action and announcement do not agree, as in Theorem 5. While the latter technique is always successful, to the extent that actions are responsive the center may be able to reduce the complexity of the mechanism by making use of the information they contain.