Informed principal problems in generalized private values environments*

Tymofiy Mylovanov† and Thomas Tröger‡

January 27, 2009

Abstract

We show that a solution to the problem of mechanism selection by an informed principal exists in a large class of environments with generalized private values: the agents' payoff functions are independent of the principal's type. The solution is an extension of Maskin and Tirole's (1990) strong unconstrained Pareto optimum. Our main condition for existence is that given any type profile the best possible outcome for the principal is the worst possible outcome for all agents. This condition is satisfied in most market environments. We also compute some examples of strong unconstrained Pareto optima.

*Financial Support by the German Science Foundation (DFG) through SFB/TR 15 "Governance and the Efficiency of Economic Systems" is gratefully acknowledged.
†Department of Economics, Penn State University. Email: txm41@psu.edu.
‡Department of Economics, University of Bonn. Email: ttroeger@uni-bonn.de.

1 Introduction

In the most of the mechanism design literature, the mechanism proposer (the principal) is assumed to have no private information. This allows to formulate the mechanism design problem as a maximization problem of the principal's payoff function subject to the agents' incentive constraints. In many circumstances, however, the assumption that the principal has no private information is not warranted: the cost of the goods to the seller who designs an auction can be unknown to bidders; the beliefs of a speculator offering a bet can be his private information; there can be uncertainty about the valuation of a supplier suggesting a collusive agreement to her competitors; a bidder who decides on a mechanism to resell some of the goods acquired in an auction can have private information about her valuations.

If the principal has private information, the proposal of a mechanism must be viewed as a move in a game and the maximization approach is not applicable. In particular, the agents posterior belief about the principals type after a mechanism is proposed may differ from their prior belief. The seminal references for this problem of informed principal are Myerson (1983) and Maskin and Tirole (1990, 1992). Nevertheless, their models exclude the standard market environments where private goods are traded.

In this paper, we provide a solution of the informed principal problem in a large class of market environments.[1] We consider environments with generalized private values: the agents' payoffs and types are independent of the principal's type, while the principal's payoff may depend on the agents' types, and the agents' payoffs are allowed to be interdependent. We permit arbitrary continuous payoff functions; in particular, we do not require a single-crossing property or risk-neutrality. Moreover, we do not impose any conditions on the outcome space, such as, e.g., the number of units of good to be allocated by the principal. Finally, we do not restrict the number of agents and their types.

Our paper builds on the work of Maskin and Tirole (1990). Maskin and Tirole discovered the crucial role played by direct mechanisms which they call Strong Unconstrained Pareto Optima (SUPO). A direct mechanism M is an SUPO if (i) M is incentive-feasible given the prior belief, and (ii) there exists no direct mechanism M' and no posterior belief G about the principal's

[1] Our results can be used to compute the solution of the informed principal problem, for example, in multi-unit auction environments with risk-neutral or even risk-averse bidders and a privately informed seller.

type such that M' is incentive-feasible given the belief G and the principal prefers M' to M independently of her type. In a private-value environment, Maskin and Tirole prove existence of SUPO and demonstrate that SUPO is the unique perfect Bayesian equilibrium outcome of an informed principal game.

To obtain their results, Maskin and Tirole envision a competitive equilibrium in a fictitious economy where the various types of the principal trade amounts of slack allowed for one specific incentive constraint and one specific participation constraint of the agent, with initial endowments being 0. Building on the arguments in Debreu (1959), they show that slack exchange equilibrium exists. Furthermore, a slack exchange equilibrium has welfare properties that ensure that it is an SUPO.

What is the main difficulty with extending this approach to environments other than the ones considered in Maskin and Tirole? The assumptions imposed by Maskin and Tirole[2] guarantee that the equilibrium prices for the two traded slacks are non-zero, and that there are no gains from trading slacks of any of the other constraints. These properties are important in establishing the connection between slack exchange equilibrium and SUPO; they do not hold in general if the assumptions of Maskin and Tirole are not satisfied.[3]

To overcome this difficulty and establish a *generally* valid connection between slack exchange equilibrium and SUPO, we allow trade in all constraints. Our crucial insight is to consider a weak version of competitive equilibrium where (i) the price of slack may be equal to 0 for some constraints and (ii) market clearing may fail for such constraints (that is, the aggregate consumption may be strictly negative).

Our main result is that a (generalized) slack exchange equilibrium exists under rather weak conditions (Proposition 3). The essential condition is that the payoff functions are such that, given any type profile, the best outcome for the principal is the worst outcome for the agents; if this condition is true,

[2]These assumptions are: (i) one agent with two types, (ii) two-dimensional outcomes, (iii) concavity, differentiability, and monotonicity of payoff functions, (iv) type-independent agent's reservation utility, (v) a condition on transfers, (vi) a condition on the reservation payoffs of the parties, and (vii) a sorting condition on the principal's payoff.

[3]For instance, in the partnership dissolution example presented in Section 2, which violates condition (vii) in Maskin and Tirole, different types of the principal consume slacks of different constraints, only one of the consumed slacks is traded, and the equilibrium prices are zero for the slack that *is* traded and non-zero for the slacks that are *not* traded.

Maskin and Tirole's assumptions can be relaxed. This best-worst-condition is satisfied in most market environments where private goods are traded.

Our slack exchange equilibrium existence proof is along the lines Maskin and Tirole (1986, 1990), who in turn follow Debreu (1959). The details of the proof are, however, significantly more demanding; for example, the payoffs of the traders in the slack exchange economy may fail to be smooth functions of the consumed vector of slack.

We also show that any slack exchange equilibrium is an SUPO (Proposition 2) and that any SUPO is a perfect Bayesian equilibrium outcome in an appropriately defined non-cooperative informed principal game; in particular, it is an equilibrium outcome both in Myerson's (1983) game and Maskin and Tirole's (1990) game (Proposition 1).

In addition, we show that the Lagrange multiplier technique proposed by Maskin and Tirole as a shortcut towards computing slack exchange equilibria applies in our framework (Lemma 3). We use this technique to establish necessary and sufficient conditions under which the equilibrium outcome in the informed principal game can be described by the mechanisms that would be optimal if the principal's information were public (Proposition 4).

Finally, we compute slack exchange equilibria in some environments, whose special cases include a discrete-type version of the Myerson and Satterthwaite bargaining environment (Myerson and Satterthwaite 1983), a version of the Akerlof's Lemons market (Akerlof 1970), a speculative trade environment with non common priors, and a partnership dissolution problem (Proposition 5 and Remark 1).

Our results are obtained for finite type spaces. Extending slack exchange equilibrium approach to continuous type spaces appears to be technically challenging as it requires considering trade in a continuum of goods.

In another paper, Mylovanov and Troeger (2008), we consider the informed principal problem in Myerson (1981) optimal auction environments and in Guesnerie and Laffont (1984) principal agent environments. These papers differ in three aspects. First, in the other paper, we show that the privacy of the principal's information does not affect the outcome of the mechanism design problem; this is not so in general environments.[4] Second, in this paper we employ Maskin and Tirole's solution concept, SUPO, whereas in the other paper we apply Myerson's strong solution. Finally, the results

[4] Maskin and Tirole (1992) demonstrate this in a generic class of environments. See also the solution of the informed principal problem computed in Remark 1.

in this paper are obtained for finite type spaces whereas all applications in the other paper are with continuous type spaces.

The rest of the paper is organized as follows. Section 2 presents the example. The model is described in Section 3. We introduce SUPO and prove that any SUPO is a perfect Bayesian equilibrium outcome of an informed principal game in Section 4. The existence of slack exchange equilibrium and SUPO is demonstrated in Section 5. We compute examples of SUPO in Section 6.

2 Example

In this section, we describe the solution of the informed principal problem, SUPO, in the following partnership dissolution environment. There are a principal (player 0) and an agent (player 1). Each of them owns a half (one share) of a company.

Let $y \in [-1, 1]$ denote the amount of shares transferred from the principal to the agent and $p \in \mathbb{R}$ denote the payment from the agent to the principal. The parties' preferences are expressed by linear risk-neutral payoff functions:

$$\begin{aligned} u_0(y, p, t_0) &= p - yt_0, \\ u_1(y, p, t_1) &= yt_1 - p, \end{aligned}$$

where t_0 and t_1 are the parties' marginal valuation of the shares (their types). Let $z_0 = (y_0, p_0) = (0, 0)$ denote the no trade outcome.

The players' types are their private information. We assume that $t_0 \in \{0, 3\}$ and $t_1 \in \{1, 2\}$. The principal believes that both agent types are equally likely. The agent believes that $t_0 = 0$ with probability $\alpha \in [0, 1]$.

The objective of the principal is to design a trading mechanism that maximizes her expected payoff subject to the individual rationality constraint that the agent agrees to participate in the mechanism. Let v_0 and v_3 denote the principal's continuation payoff if $t_0 = 0$ and $t_0 = 3$ respectively.

If the principal's type were common knowledge, the principal would obtain the maximal feasible payoff of $v_0 = v_3 = 1$ by offering, for example, to sell at the price of 1 if her type is low and buy at the price of 2 is her type is high.

Let V_α denote the set of principal's payoffs that can be obtained in some individually rational and incentive compatible direct mechanism and are higher than what the principal can obtain when her type is commonly known. In our model,

$$V_\alpha = \{(v_0, v_3) | \alpha v_0 + (1-\alpha) v_3 \leq 2 - \alpha, v_0 \geq 1, v_3 \geq 1\}.$$

In Section 4, we introduce the notion of mechanisms that are Strong Unconstrained Pareto Optimum (SUPO). These mechanisms are characterized by the Pareto property that for any prior beliefs about the principal there does not exist another incentive feasible mechanism in which all types of the principal are better off and some are strictly better off than in SUPO. In our example, we can describe the set of the principal's payoffs attainable in SUPO as the payoffs that belong to V_α and are not dominated by any payoffs in $V_{\alpha'}$ for any $\alpha' \in [0,1]$,

$$W_\alpha = \{(v_0, v_3) \in V_\alpha | v_0 \geq v_0' \text{ and } v_3 \geq v_3' \text{ for any } \alpha' \in [0,1] \text{ and } (v_0', v_3') \in V_{\alpha'}\}$$

It is straightforward to see that this set is not empty and is given by

$$W_\alpha = \begin{cases} \{(1,2)\}, & \text{if } \alpha < 1/2; \\ V_{\frac{1}{2}}, & \text{if } \alpha = 1/2; \\ \{(2,1)\}, & \text{if } \alpha > 1/2. \end{cases}$$

We offer an example of a simple indirect mechanism that implements the payoffs in W_α: there is a fixed price p and the principal has the right to choose whether to buy or sell at this price. The price is equal to 1 if $\alpha > 1/2$, 2 if $\alpha < 1/2$, and any value in $[1,2]$ if $\alpha = 1/2$.

3 Model

We consider the interaction of a principal (player 0) and n agents (players $i \in N = \{1, \ldots, n\}$). The players must collectively choose an outcome from a measurable space of *basic outcomes* Z. Every player $i = 0, \ldots, n$ has a *type* t_i that belongs to a finite *type space* T_i. The product of agents' type spaces is denoted $\mathbf{T} = T_0 \times \cdots \times T_n$. Player i's payoff function is denoted

$$u_i : \; Z \times \mathbf{T} \to \mathbb{R},$$

That is, player i's payoff can depend on the outcome and on every player's type.

We assume that $u_i(\cdot,\mathbf{t}): Z \to \mathbb{R}$ is measurable for all $\mathbf{t} \in \mathbf{T}$, and that u_i is bounded. The types $t_0,\ldots,t_n$ are realizations of stochastically independent random variables with probability distributions $P_0,\ldots,P_n$, where $P_i = (p_i^{(t_i)})_{t_i \in T_i}$. We call P_i the *prior distribution* for player i's type. The joint distribution of players' types is denoted $\mathbf{P}$. We will use the notation $\mathbf{t}_{-i}$ for the vector of types of the players other than i, use $\mathbf{T}_{-i}$ for the respective product of type spaces, and use $\mathbf{P}_{-i}$ for the respective product of distributions. Similarly, we use the index $-i-0$ if agent i as well as the principal are excluded.

The interaction leads to a probability distribution over basic outcomes; let $\mathcal{Z}$ denote the set of probability measures on Z. Any element of $\mathcal{Z}$ is called an *outcome*. We endow $\mathcal{Z}$ with the smallest σ-algebra such that, for every measurable set $B \subseteq Z$, the mapping $m_B: \ \mathcal{Z} \to [0,1], \ \zeta \mapsto \zeta(B)$ is measurable.[5] We identify any $z \in Z$ with the point distribution that puts probability 1 on the point z; hence, $Z \subseteq \mathcal{Z}$.[6] We extend the definition of u_i to $\mathcal{Z} \times \mathbf{T}$ via the statistical expectation:[7]

$$u_i(\zeta,\mathbf{t}) \;=\; \int_Z u_i(z,\mathbf{t})\zeta(\mathrm{d}z).$$

Some outcome $z_0 \in \mathcal{Z}$ is designated as the *disagreement outcome.*

The interaction is described by the following *informed-principal game.* First, each player privately observes her type t_i. Second, the principal offers a mechanism M (a precise definition is given later). Third, the agents decide simultaneously whether or not to accept M. If M is accepted unanimously, each player chooses a message in M, and the outcome specified by M is implemented. If at least one agent rejects M, the disagreement outcome z_0

[5]Given this σ-algebra, any uncertainty about outcomes in $\mathcal{Z}$ can be equivalently described as uncertainty about basic outcomes in Z. Formally, any probability measure P on $\mathcal{Z}$ can be identified with a probability measure ζ_P on Z, via the definition

$$\zeta_P(B) \;=\; \int_{\mathcal{Z}} \zeta(B)P(\mathrm{d}\zeta) \quad \text{for every measurable } \ B \subseteq Z.$$

[6]Observe that, if $\mathcal{M}$ is an arbitrary measurable space and if a mapping $f: \mathcal{M} \to Z$ is measurable with respect to the σ-algebra on Z, then f is also measurable when viewed as a mapping into $\mathcal{Z}$ (the reason is that the composite mapping $m_B f$ is measurable for every measurable $B \subseteq Z$).

[7]Observe that the extended mapping $u_i: \ \mathcal{Z} \times \mathbf{T} \to \mathbb{R}$ inherits the following properties: the function $u_i(\cdot,\mathbf{t}): \mathcal{Z} \to \mathbb{R}$ is measurable for all $\mathbf{t} \in \mathbf{T}$ and u_i is bounded.

is implemented.[8]

An *allocation rule* is a function

$$\rho: \quad \mathbf{T} \to \mathcal{Z}, \quad \mathbf{t} \mapsto \rho(\mathbf{t})$$

that assigns an outcome $\rho(\mathbf{t})$ to every type profile $\mathbf{t}$. Thus, an allocation rule describes the outcome of the players' interaction as a function of the type profile. Alternatively, an allocation rule ρ can be interpreted as a *direct mechanism*, where the players $i = 0, \ldots, n$ simultaneously announce types $\hat{t}_i$ (=messages), and the outcome $\rho(\hat{t}_0, \ldots, \hat{t}_n)$ is implemented.

Let Q_0 denote the probability distribution that describes the agents' belief about the principal's type if the direct mechanism ρ is accepted. The expected payoff of type t_i of player i if she announces the type $\hat{t}_i$ while all other players announce their types truthfully, is

$$U_i^{\rho,Q_0}(\hat{t}_i, t_i) \;=\; \sum_{\mathbf{T}_{-i}} u_i(\rho(\hat{t}_i, \mathbf{t}_{-i}), (t_i, \mathbf{t}_{-i}))\; \mathbf{q}_{-i}(\mathbf{t}_{-i}),$$

where $\mathbf{q}_{-i} = q_0 \times \mathbf{p}_{-i-0}$ if $i \neq 0$, and $\mathbf{q}_{-0} = \mathbf{p}_{-0}$. We will use the shortcuts $U_i^{\rho,Q_0}(t_i) = U_i^{\rho,Q_0}(t_i, t_i)$ and $U_0^{\rho} = U_0^{\rho,Q_0}$. The expected payoff if the mechanism ρ is rejected is denoted[9]

$$\underline{U}_i^{Q_0}(t_i) \;=\; \sum_{\mathbf{T}_{-i}} u_i(z_0, (t_i, \mathbf{t}_{-i}))\; \mathbf{q}_{-i}(\mathbf{t}_{-i}).$$

A direct mechanism ρ is called Q_0-*incentive feasible* if no type of any player has an incentive to deviate from announcing her true type or can gain from refusing to participate: for all i, t_i, $\hat{t}_i$,

$$U_i^{\rho,Q_0}(t_i) \;\geq\; U_i^{\rho,Q_0}(\hat{t}_i, t_i), \tag{1}$$
$$U_i^{\rho,Q_0}(t_i) \;\geq\; \underline{U}_i^{Q_0}(t_i). \tag{2}$$

[8]This game differs slightly from the games specified by Myerson (1983) and by Maskin and Tirole (1990, 1992). Myerson allows the possibility of other private actions beyond acceptance and rejection, and assumes that private actions and messages in M are chosen simultaneously. Maskin and Tirole assume that players can use a public randomization device to decide which equilibrium to play in M.

[9]Observe that when computing her expected payoff, agent i uses the prior beliefs about the other agents; this is appropriate if she expects all other agents to accept and if she actually is the only one to reject.

An P_0-incentive feasible allocation rule is simply called *incentive feasible*. A direct mechanism ρ is called Q_0-*unconstrained feasible* if (1) holds for all $i \neq 0$, and (2) holds for all i; i.e., the definition of unconstrained feasibility ignores the possibility that the principal may have an incentive to deviate from announcing her true type.

Generalized private values environments

Our results will mainly concern environments with *generalized private values*, where the agents' payoff functions are independent of the principal's type, that is, for all $i \geq 1$,

$$u_i(z, (t_0, \mathbf{t}_{-0})) = u_i(z, (t'_0, \mathbf{t}_{-0})) \quad \text{for all } \; z, \; t_0, t'_0, \; \mathbf{t}_{-0}.$$

Many interesting environments belong to this class. All private-environments are obvious members. But the principal's payoff can also depend on the agents' types, and the agents can have interdependent types. For example, a variant of a Lemons market where the seller (agent) is, as usual, privately informed about quality, and the buyer (principal) is privately informed about marginal willingness-to-pay for additional quality, has generalized private values.

4 Strong Unconstrained Pareto Optimum

Generalizing Maskin and Tirole (1990), we introduce *strong unconstrained Pareto optimum (SUPO)*. This concept turns out to be a very useful solution concept for a large class of informed-principal problems.

For any two allocation rules ρ and ρ', let the sets of types of the principal that are strictly better off in either allocation rule be denoted

$$\begin{aligned} S_>(\rho, \rho') &= \{t_0 \in T_0 \mid U_0^{\rho'}(t_0) > U_0^{\rho}(t_0)\}, \\ S_<(\rho, \rho') &= \{t_0 \in T_0 \mid U_0^{\rho'}(t_0) < U_0^{\rho}(t_0)\}. \end{aligned}$$

An allocation rule ρ is *unconstrained-dominated* by an allocation rule ρ' if there exists a belief Q_0 such that ρ' is Q_0-unconstrained feasible and

$$\begin{aligned} \Pr_{Q_0}(S_>(\rho, \rho')) &> 0, \\ S_<(\rho, \rho') &= \emptyset. \end{aligned}$$

An incentive feasible allocation rule that is not unconstrained dominated is a *strong unconstrained Pareto optimum* (SUPO, Maskin and Tirole, (1990)).[10] SUPO appears to be a very restrictive concept. However, as we will see, an SUPO exists in a large class of generalized private values environments.

One should distinguish the unconstrained-domination concept from Myerson's (1983) stronger concept of domination. An incentive feasible allocation rule ρ is *dominated* if there exists an incentive feasible allocation rule ρ' such that all types of the principal are at least as well off in ρ' as in ρ, and a F_0-positive mass of types of the principal is strictly better off in ρ'.[11] Clearly, if an allocation is dominated, then it is unconstrained dominated.

Perfect Bayesian equilibrium

Any SUPO is a perfect Bayesian equilibrium in an appropriately specified informed-principal game. Specifying the game is, unfortunately, not straightforward because a careful definition needs to be made about what constitutes a "mechanism". As observed by Myerson (1983) and Maskin and Tirole (1990), the standard approach of applying the revelation principle and restricting attention to direct mechanisms is not possible. The principal's very act of offering a particular mechanism may force the agents to update their belief about the principal's type, away from the prior F_0. Hence, whatever mechanism is proposed, one must consider its equilibria for all possible beliefs about the principal.[12]

On the equilibrium path, however, we can assume, without loss of generality, that the same direct mechanism is proposed by all types of the principal

[10] Formally, our definition differs slightly from Maskin and Tirole's (1990). However, if both Maskin and Tirole's and our assumptions hold, then our definition is equivalent to theirs. This can be seen from the proofs of Propositions 3 and 4 in Maskin and Tirole's paper.

[11] Even if one restricts attention to generalized private values environments, SUPO differs from the solution concepts *neutral optimum* and *strong solution* proposed by Myerson (1983): neutral optimum may exist when no SUPO exists, and an SUPO may exist when no strong solution exists. However, if both a strong solution and an SUPO exist, then they lead to identical payoffs for all types of the principal (this holds in general environments).

[12] An interesting example by Yilankaya (1999) shows that, in general, the agents' off-path beliefs will have to be different from the prior beliefs in order to support an equilibrium. His example involves a bilateral trade environment, where ρ is constructed from optimal fixed-price offers by all types of the seller (principal), and M is a double auction. If the agent believes in the lowest-cost type of the seller, then all types can be prevented from deviating to the double auction.

and is accepted by all agents (Myerson's "inscrutability principle"). This follows from the revelation principle: any perfect Bayesian equilibrium of the informed-principal game induces some incentive feasible allocation ρ; without loss of generality, all types of the principal offer the direct mechanism ρ. Perfect Bayesian equilibrium in the informed principal game then requires that no type of the principal has an incentive to deviate by offering a mechanism $M \neq \rho$, given that the continuation play is sequentially optimal given the agents' (off-path) belief about the principal. As usual, off-path beliefs can be arbitrarily defined.

A *finite mechanism* is a finite multi-stage game form with observed actions, with players $N \cup \{0\}$, and with outcomes $\mathcal{Z}$.[13] (Observe that the informed-principal game where any finite mechanism is feasible is *not* a finite game because the set of finite mechanisms is infinite.)

Proposition 1. *Let ρ be an SUPO. Consider the informed principal game where the set of feasible mechanisms consists of all finite mechanisms.*

Then there exists a perfect Bayesian equilibrium where all types of the principal propose the direct mechanism ρ, and the allocation rule ρ is the equilibrium outcome.

Proof. Define a function w on T_0 by $w(t_0) = U^\rho(t_0)$.

For any regular mechanism $M \neq \rho$, consider the following reduced game (M, w):

> At stage 1, each type t_0 of the principal chooses between obtaining the payoff $w(t_0)$, which ends the game (and the agents' payoffs are irrelevant), or proposing the mechanism M. If M is proposed, then the agents decide about acceptance at stage 2. If M is unanimously accepted, then M is played at stage 3, otherwise the disagreement outcome z_0 is implemented.

This is finite multi-stage game with observable actions and thus has a perfect Bayesian equilibrium, denoted σ_M (cf., e.g., Fudenberg and Tirole, Section 8.2.3.).

[13] It seems appropriate to consider the possibility that the principal proposes a multi-stage mechanism rather than the strategic form of this mechanism, in order to reduce the set of possible (perfect Bayesian) equilibrium outcomes of the mechanism, thus influencing the players' actions in the mechanism. In the earlier literature (Myerson, (1983), and Maskin and Tirole, (1990)) only normal-form mechanisms are considered.

Let ρ_M denote the allocation rule induced by the continuation equilibrium at the beginning of stage 2 in (M, w), and let Q_M denote the belief about the principal at the beginning of stage 2 in (M, w).

We construct a PBE of the informed-principal game as follows. The principal proposes the direct mechanism ρ, everybody accepts, and everybody announces their true type. If the principal proposes a mechanism $M \neq \rho$, then the agents hold the belief Q_M. Moreover, if M is proposed, then the continuation strategies are as in σ_M.

It remains to be shown that no type of the principal has an incentive to deviate from proposing ρ to proposing any $M \neq \rho$; i.e., to show that

$$S_>(\rho, \rho_M) \quad = \quad \emptyset. \tag{3}$$

Let ρ' be the allocation that coincides with ρ_M for all principal-types in the support of Q_M, and otherwise coincides with ρ. If (3) does not hold, then ρ is unconstrained dominated by the Q_M-incentive feasible allocation rule ρ', a contradiction to the assumption that ρ is an SUPO. *QED*

Proposition 1 generalizes a result of Maskin and Tirole (1990). Our proof is different from theirs. In particular, we do not rely on the connection between SUPO and Walrasian equilibrium.

The logic of the proof of Proposition 1 extends straightforwardly to environments with a non-finite type space. However, the definition of a "mechanism" then needs to include infinite game forms (even direct mechanisms are then infinite game forms), and some technical restrictions need to be added to avoid equilibrium non-existence.

5 Existence of SUPO in generalized private values environments

In this section, we assume that the basic outcome space Z is a compact metric space such that all payoff functions u_i ($i \geq 0$) are continuous. Similar to Maskin and Tirole (1990), we define an exchange economy where the different types of the principal trade amounts of slack granted to the incentive constraints and participation constraints of the agents. The crucial departure from Maskin and Tirole's approach is that we consider trade in *all* constraints.

For all t_0 and real-valued functions r on

$$\mathcal{R} \;=\; \cup_{i\geq 1}\{i\} \times T_i$$

and c on

$$\mathcal{C} \;=\; \cup_{i\geq 1}\{i\} \times \{(\hat{t}_i, t_i) \mid \hat{t}_i, t_i \in T_i,\; \hat{t}_i \neq t_i\},$$

consider the problem

$$\begin{aligned}
P(t_0, r, c): \qquad & \max_{\rho:\mathbf{T}\to\mathcal{Z}} \sum_{\mathbf{t}_{-0}} u_0(\rho(\mathbf{t}), \mathbf{t})\mathbf{p}_{-0}(\mathbf{t}_{-0}) \\
\text{s.t.} \quad & \sum_{\mathbf{t}_{-0-i}} (u_i(\rho(\mathbf{t}), \mathbf{t}) - u_i(z_0, \mathbf{t}))\mathbf{p}_{-0-i}(\mathbf{t}_{-0-i}) \geq -r(i, t_i) \\
& \qquad \text{for all } (i, t_i) \in \mathcal{R}, \\
& \sum_{\mathbf{t}_{-0-i}} (u_i(\rho(\mathbf{t}), \mathbf{t}) - u_i(\rho(\hat{t}_i, \mathbf{t}_{-i}), \mathbf{t}))\mathbf{p}_{-0-i}(\mathbf{t}_{-0-i}) \geq -c(i, \hat{t}_i, t_i) \\
& \qquad \text{for all } (i, \hat{t}_i, t_i) \in \mathcal{C}.
\end{aligned}$$

According to the problem $P(t_0, r, c)$, type t_0 of the principal maximizes her expected payoff, given certain (positive or negative) slacks in the agents' constraints, as described by the functions r and c.

Let C denote the set of (r, c) such that the constraint set of problem $P(t_0, r, c)$ is non-empty and such that (r, c) is bounded. Observe that C is non-empty (the point where both r and c are identically 0 belongs to C because the allocation rule that implements the disagreement outcome satisfies all constraints). Moreover, C is convex.

Lemma 1. *Problem $P(t_0, r, c)$ has a solution, for all $t_0 \in T_0$ and all $(r, c) \in C$.*

Proof. When we endow $\mathcal{Z}$ with the weak topology, then, for any given $\mathbf{t}$ the functions $u_0(\cdot, \mathbf{t})$ and $u_i(\cdot, \mathbf{t})$ are continuous as functions of $\mathcal{Z}$, and $\mathcal{Z}$ is a compact metric space, by Prohorov's Theorem. Hence, with respect to the product topology on $\mathcal{Z}^{|\mathbf{T}|}$, the objective of $P(t_0, r, c)$ is continuous and the constraint set is compact. Hence, a maximizer exists. *QED*

For all $(r, c) \in C$, let $V(t_0, r, c) < \infty$ denote the maximum value of problem $P(t_0, r, c)$.

For all $i \geq 1$, endow $\{i\} \times T_i$ with the same topology as T_i. Endow $\mathcal{R}$ with the standard induced topology for disjoint unions.

Let β be a non-negative function on $\mathcal{R}$, and γ be a non-negative function on $\mathcal{C}$ such that

$$\sum_{i\geq 1,\ t_i} \beta(i,t_i) \ + \sum_{i\geq 1,\ \hat{t}_i \neq t_i} \gamma(i,\hat{t}_i,t_i) \ > \ 0.$$

We write (β,γ) for the induced function on $\mathcal{R} \cup \mathcal{C}$.

For any (β,γ) and any functions r on $\mathcal{R}$ and c on $\mathcal{C}$, let

$$\begin{aligned}\beta \cdot r &= \sum_{i\geq 1,\ t_i} \beta(i,t_i) r(i,t_i),\\ \gamma \cdot c &= \sum_{i\geq 1,\ \hat{t}_i \neq t_i} \gamma(i,\hat{t}_i,t_i) c(i,\hat{t}_i,t_i).\end{aligned}$$

A *slack exchange equilibrium* is a list

$$(r^*_{t_0}, c^*_{t_0})_{t_0 \in T_0}, \beta^*, \gamma^*$$

such that (β^*,γ^*) is non-negative and not identically zero on $\mathcal{R} \cup \mathcal{C}$, and, for all $t_0 \in T_0$,

$$V(t_0, r^*_{t_0}, c^*_{t_0}) = \max_{(r,c)\in C} V(t_0,r,c) \quad \text{s.t.} \ \ \beta^* \cdot r + \gamma^* \cdot c \leq 0, \tag{4}$$

$$\beta^* \cdot r + \gamma^* \cdot c = 0 \ \text{ for all maximizers } (r,c) \text{ in (4)}, \tag{5}$$

and, for all $i \geq 1$, t_i, and $\hat{t}_i \neq t_i$,

$$\sum_{T_0} r^*_{t_0}(i,t_i)\ p_0(t_0) \ \leq \ 0, \tag{6}$$

$$\sum_{T_0} c^*_{t_0}(i,\hat{t}_i,t_i)\ p_0(t_0) \ \leq \ 0. \tag{7}$$

Each type t_0 of the principal can be interpreted as a trader in an exchange economy. Slacking a constraint (i,t_i) by an amount $r(i,t_i)$ can be interpreted as consuming the (positive or negative) quantity $r(i,t_i)$ of a good (i,t_i). Similarly, there are goods $(i,\hat{t}_i,t_i)$. Each trader t_0 has an initial endowment of 0 of each good.

Given the price vector (β^*, γ^*), each trader t_0 optimally decides which bundle of slacks $(r^*_{t_0}, c^*_{t_0})$ to buy (4). Walras' law holds (5). The total consumption of each good does not exceed the total initial endowment (6, 7).[14]

Proposition 2. *Any slack exchange equilibrium in a generalized private values environment is an SUPO.*

Proof. Let ρ be a maximizer of problem $P(t_0, r^*_{t_0}, c^*_{t_0})$ for all t_0. By (6) and (7), ρ satisfies constraints (2) and (1) for all $i \geq 1$. Because the allocation rule that implements the disagreement outcome is feasible in problem $P(t_0, 0, 0)$, (2) is satisfied for $i = 0$. Because, for all $t_0, \hat{t}_0$, the bundle $(r^*_{\hat{t}_0}, c^*_{\hat{t}_0})$ belongs to the constraint set of problem (4), constraint (1) is satisfied for $i = 0$. In summary, ρ is incentive feasible.

To complete the proof ρ is an SUPO, suppose that ρ is unconstrained-dominated by an allocation rule ρ'; let Q_0 denote the corresponding belief. Then, there exists a set S with positive Q_0-measure such that

$$\begin{aligned} & U_0^{\rho'}(t_0) \geq U_0^{\rho}(t_0) \quad \text{for all } t_0 \in T_0, \\ \text{and} \quad & U_0^{\rho'}(t'_0) > U_0^{\rho}(t'_0) \quad \text{for all } t'_0 \in S. \end{aligned} \tag{8}$$

For all t_0, define r'_{t_0} and c'_{t_0} such that ρ' satisfies all constraints of problem $P(t_0, r'_{t_0}, c'_{t_0})$ with equality. Because ρ' is unconstrained-feasible,

$$\sum_{T_0} r'_{t_0}(i, t_i)\; q_0(t_0) \;\leq\; 0 \quad \text{for all } \; i, t_i, \tag{9}$$

$$\sum_{T_0} c'_{t_0}(i, \hat{t}_i, t_i)\; q_0(t_0) \;\leq\; 0 \quad \text{for all } \; i, \hat{t}_i, t_i. \tag{10}$$

Because $(r^*_{t_0}, c^*_{t_0})$ satisfies (4) and (5), (8) implies

$$\beta^* \cdot r'_{t_0} + \gamma^* \cdot c'_{t_0} \geq 0, \quad \text{and ``>'' if } t_0 \in S.$$

Adding over T_0 according to the measure Q_0, we obtain a contradiction to (9) or (10). *QED*

We consider trade in *all* constraints, whereas Maskin and Tirole consider trade in just two constraints in a specific class of economic environments with

[14] We allow that part of the initial endowment is destroyed. This can be relevant for goods with price 0 and is potentially important for equilibrium existence.

one agent and two types, where these two constraints cannot be relaxed and the other two constraints are automatically satisfied. Because we consider trade in all constraints, the equilibrium prices of some constraints may be 0. For any such constraint, the aggregate amount of slack "consumed" in equilibrium may be strictly below the aggregate endowment of 0. All of these differences could be relevant: in Section 6, we present an example of a speculative trade environment in which (i) all constraints could be potentially relevant, (ii) only one constraint is traded at 0 price, and (iii) the aggregate consumption of the slack of the constraint which is traded is negative.

The focus on generalized private values environments is essential to make the exchange-economy-technique applicable: it guarantees that the form of the agents' incentive compatibility and participation constraints is independent of the type of the principal.

We say that *any best outcome for the principal is a worst outcome for all agents* if

$$\arg\max_{z\in Z} u_0(z,\mathbf{t}) \subseteq \arg\min_{z\in Z} u_i(z,\mathbf{t}) \quad \text{for all} \quad \mathbf{t}.$$

This condition is useful because it guarantees Walras' law for the fictitious economy.[15]

Lemma 2. *If any best outcome for the principal is a worst outcome for all agents, then (5) holds for all t_0, β^*, and γ^*.*

Proof. Suppose that ρ is a maximizer of problem $P(t_0,r,c)$ and $\beta^*\cdot r+\gamma^*\cdot c<0$. Then there exists $i\geq 1$ and t_i such that $r(i,t_i)<0$ or $i\geq 1$, $\hat{t}_i$ and t_i such that $c(i,\hat{t}_i,t_i)<0$. Hence, looking at the constraints of problem $P(t_0,r,c)$, there exists i, t_i such that $U_i^\rho(t_i)$ is not the lowest feasible expected payoff for type t_i of agent i. Hence, there exists a type profile t_{-0-i} such that $\rho(\mathbf{t})$ puts probability less than 1 on the outcomes in $\arg\min_{z\in Z} u_i(z,\mathbf{t})$. Because any best outcome for the principal is a worst outcome for all agents, $\rho(\mathbf{t})$ puts probability less than 1 on the outcomes in $\arg\max_{z\in Z} u_0(z,\mathbf{t})$.

Consider $(r',c'):=(r+\epsilon,c+\epsilon)$ with $\epsilon>0$ so small that

$$\beta\cdot r'+\gamma\cdot c' \;<\; 0. \tag{11}$$

The allocation ρ satisfies all constraints of problem $P(t_0,r',c')$ with strict inequality. Let ρ' be an allocation rule that, for the type profile $\mathbf{t}$ constructed

[15] I.e., the condition replaces the condition of locally non-satiated preferences that is needed to guarantee the existence of a Walrasian equilibrium in the standard context.

above, implements any outcome in $\arg\max_{z\in Z} u_0(z, \mathbf{t})$, and for all other type profiles implements the same outcome as ρ. Then, an allocation rule ρ'' that implements ρ with probability $\lambda < 1$ and ρ' with probability $1-\lambda$, belongs to the constraint set of problem $P(t_0, r', c')$ if λ is sufficiently close to 1, and yields a higher value for the objective of $P(t_0, r', c')$ than ρ. Hence, $V(t_0, r', c') > V(t_0, r, c)$. Hence, (r, c) is not a maximizer of the problem in (4), because, by (11), the point (r', c') satisfies the constraint of this problem. QED

An environment with generalized private values is *constraint-non-degenerate* if there exists an allocation rule such that, for any $Q_0 = 1_{t_0}$,[16] the incentive constraints (1) and participation constraints (2) are satisfied with strict inequality for all agents $i \geq 1$, $\hat{t}_i$, and t_i.

Proposition 3. *Consider any generalized private values constraint-non-degenerate environment, where any best outcome for the principal is a worst outcome for all agents.*

Then a slack exchange equilibrium exists.

Before we prove this result, comments are in order. Our basic line of proof is analogous to Maskin and Tirole (1986, 1990), who in turn follow Debreu (1959). The proof is technically more demanding than Maskin and Tirole's. In particular, we cannot show that the traders' utility functions in the fictitious economy are continuous at the boundary of the consumption set.

The assumption of finite type spaces greatly simplifies the technicalities because it guarantees that the fictitious economy has finitely many traders and finitely many goods. The assumption of constraint non-degeneracy is essential towards showing that the demand correspondence in the fictitious economy is upper hemicontinuous. The assumption that any best outcome for the principal is a worst outcome for all agents guarantees that Walras' law holds (cf. Lemma 2).

Proof of Proposition 3. Observe that the set C is closed. (Consider any sequence $(r^m, c^m) \to (r, c)$ such that $(r^m, c^m) \in C$. By assumption, the constraint set of $P(t_0, r^m, c^m)$ contains a point ρ^m. For all sufficiently large m, the point ρ^m belongs to the constraint set of problem $P(t_0, r+1, c+1)$, where 1 denote the function that is identically equal to 1. Because the latter

[16]Due to generalized private values, it is irrelevant which t_0 is used here.

constraint set is compact, ρ^m has a subsequence that converges to some point ρ'. By continuity, ρ' belongs to the constraint set of $P(t_0, r, c)$. Hence, $(r, c) \in C$.)

Because Z is compact, there exists an upper bound for the size of the left hand side of every constraint of $P(t_0, r, c)$. Hence, there exist functions $\overline{r}$ on $N \times T_i$ and $\overline{c}$ on $N \times T_i^2$ such that

$$V(t_0, r, c) \;=\; V(t_0, \min\{r, \overline{r}\}, \min\{c, \overline{c}\}) \quad \text{for all} \;\; (r, c) \in C, \tag{12}$$

where "min" determines the point-wise minimum of two functions.

Similarly, there exist functions $\underline{r}$ on $N \times T_i$ and $\underline{c}$ on $N \times T_i^2$ such that

$$C \;\subseteq \{(r, c) \mid r \geq \underline{r},\; c \geq \underline{c}\}, \tag{13}$$

where "$\geq$" refers to point-wise comparison. By (12) and (13), the set

$$D \;=\; C \cap \{(r, c) \mid r \leq \overline{r},\; c \leq \overline{c}\} \quad \text{is compact,} \tag{14}$$

where "$\leq$" refers to point-wise comparison.

Define the unit simplex

$$\begin{aligned} \Delta \;=\; \{(\beta, \gamma) \mid\ & \beta : \cup_{i\geq 1}\{i\} \times T_i \to \mathbb{R},\; \gamma : \cup_{i\geq 1}\{i\} \times T_i^2 \to \mathbb{R}, \\ & \beta(i, t_i) \geq 0,\; \gamma(i, \hat{t}_i, t_i) \geq 0, \\ & \sum_{i\geq 1,\; t_i} \beta(i, t_i) + \sum_{i\geq 1,\; \hat{t}_i \neq t_i} \gamma(i, \hat{t}_i, t_i) = 1\}. \end{aligned}$$

For all $(\beta, \gamma) \in \Delta$, consider the problem

$$E(t_0, \beta, \gamma) : \qquad \max_{(r,c)\in D} V(t_0, r, c) \quad \text{s.t.} \;\; \beta \cdot r + \gamma \cdot c \leq 0.$$

The objective $V(t_0, \cdot)$ of problem $E(t_0, \beta, \gamma)$ does not "jump downwards":[17] for any convergent sequence (x^m) in D,

$$V(t_0, \lim_m x^m) \;\geq\; \limsup_m V(t_0, x^m). \tag{15}$$

[17]It is not clear whether $V(t_0, \cdot)$ is continuous. Continuity does not follow from Berge's Theorem, because it is not clear whether the constraint set of $P(t_0, r, c)$ is lower-hemicontinuous in (r, c). Although $V(t_0, \cdot)$ is concave and hence is continuous in the interior of D, continuity at the boundary is not clear.

To see (15), let $x = \lim x^m$ and let (x^{m_l}) be a subsequence such that $V(t_0, x^{m_l})$ converges. Let ρ^l be a maximizer of problem $P(t_0, x^{m_l})$ and let (ρ^{l_k}) be a subsequence such that ρ^{l_k} converges. Because $\lim_k x^{m_{l_k}} = x$, the limit $\rho' = \lim_k \rho^{l_k}$ belongs to the constraint set of problem $P(t_0, x)$. Hence,

$$V(t_0, x) \geq U_0^{\rho'}(t_0) = \lim_k U_0^{\rho^{l_k}}(t_0) = \lim_k V(t_0, x^{m_{l_k}}) = \lim_l V(t_0, x^{m_l}).$$

By (15) and because, by (14), the constraint set of problem $E(t_0, \beta, \gamma)$ is compact, a maximizer to problem $E(t_0, \beta, \gamma)$ exists; let $e(t_0, \beta, \gamma)$ denote the set of maximizers.

The correspondence $e(t_0, \cdot) : \Delta \to D$ is convex-valued (because $V(t_0, \cdot)$ is concave). To show that $e(t_0, \cdot)$ is upper-hemicontinuous, we begin by showing that, for every sequence in Δ,

$$\text{if} \quad (\beta^m, \gamma^m) \to (\beta, \gamma) \quad \text{then} \quad \liminf_m v(t_0, \beta^m, \gamma^m) \quad \geq \quad v(t_0, \beta, \gamma), \tag{16}$$

where $v(t_0, x)$ denotes the value reached at the maximum of problem $E(t_0, x)$.

Let $(r^*, c^*) \in e(t_0, \beta, \gamma)$. If $r^* \cdot \beta + c^* \cdot \gamma < 0$, then $r^* \cdot \beta^m + c^* \cdot \gamma^m < 0$ if m is sufficiently large, hence (r^*, c^*) belongs to the constraint set of $E(t_0, \beta^m, \gamma^m)$, which shows (16). Now suppose that

$$r^* \cdot \beta + c^* \cdot \gamma \quad = \quad 0. \tag{17}$$

Because of constraint-non-degeneracy, the set D contains a strictly negative point (r^-, c^-). For all large m, define

$$\alpha^m \quad = \quad \min \left\{ 1, \frac{-(r^- \cdot \beta^m + c^- \cdot \gamma^m)}{r^* \cdot \beta^m + c^* \cdot \gamma^m - (r^- \cdot \beta^m + c^- \cdot \gamma^m)} \right\}. \tag{18}$$

Using the shortcuts $x^* = (r^*, c^*)$ and $x^- = (r^-, c^-)$, the convex combination $x^m = \alpha^m x^* + (1 - \alpha^m) x^- \in D$. By construction, x^m belongs to the constraint set of problem $E(t_0, \beta^m, \gamma^m)$. Hence, using the concavity of $V(t_0, \cdot)$,

$$\alpha^m V(t_0, x^*) + (1 - \alpha^m) V(t_0, x^-) \quad \leq \quad V(t_0, x^m) \leq v(t_0, \beta^m, \gamma^m). \tag{19}$$

As $m \to \infty$, we have $\alpha^m \to 1$ by (17) and (18). Hence, (19) implies

$$V(t_0, x^*) \quad \leq \quad \liminf_m v(t_0, \beta^m, \gamma^m).$$

Because $V(t_0, x^*) = v(t_0, \beta, \gamma)$, we obtain (16).

To show that $e(t_0,\cdot)$ is upper hemi-continuous, suppose that $(\beta^m,\gamma^m)\to(\beta,\gamma)$, $x^m\in e(t_0,\beta^m,\gamma^m)$ and $x^m\to x$. Then,

$$V(t_0,x)\overset{(15)}{\geq}\liminf_m V(t_0,x^m)=\liminf_m v(t_0,\beta^m,\gamma^m)\overset{(16)}{\geq}v(t_0,\beta,\gamma).$$

Hence, $x\in e(t_0,\beta,\gamma)$ because x belongs to the constraint set of $E(t_0,\beta,\gamma)$.

Define a correspondence $h:\prod_{t_0\in T_0}D\to\Delta$ by letting $h((r_{t_0},c_{t_0})_{t_0\in T_0})$ be the set of solutions to the problem

$$R((r_{t_0},c_{t_0})_{t_0\in T_0}):\quad \max_{(\beta,\gamma)\in\Delta}\sum_{t_0\in T_0}p(t_0)\left(\beta\cdot r_{t_0}+\gamma\cdot c_{t_0}\right).$$

By Berge's Theorem, h is upper-hemicontinuous. Moreover, h is convex-valued. By Kakutani's Theorem, the correspondence

$$\begin{aligned}(\prod_{t_0\in T_0}D)\times\Delta &\rightarrow (\prod_{t_0\in T_0}D)\times\Delta,\\ (x,(\beta,\gamma)) &\mapsto (\prod_{t_0\in T_0}e(t_0,\beta,\gamma))\times h(x)\end{aligned}$$

has a fixed point $((r^*_{t_0},c^*_{t_0})_{t_0\in T_0},(\beta^*,\gamma^*))$.

Using the constraint of problem $E(t_0,\beta^*,\gamma^*)$ for all t_0,

$$\sum_{t_0\in T_0}p(t_0)\left(\beta^*\cdot r^*_{t_0}+\gamma^*\cdot c^*_{t_0}\right)\;\leq\;0.$$

Hence,

$$\sum_{t_0\in T_0}p(t_0)r^*_{t_0}(i,t_i)\;\leq\;0\quad\text{for all }\;i,t_i. \tag{20}$$

(If not, choose $(\beta,\gamma)\in\Delta$ such that $\gamma=0$ and $\beta(j,t'_i)=0$ for all $(j,t'_i)\neq(i,t_i)$, hence $\sum_{t_0\in T_0}p(t_0)(\beta\cdot r^*_{t_0}+\gamma c^*_{t_0})>0$, which contradicts the fact that (β^*,γ^*) solves problem $R((r^*_{t_0},c^*_{t_0})_{t_0\in T_0})$). Similarly,

$$\sum_{t_0\in T_0}p(t_0)c^*_{t_0}(i,\hat{t}_i,t_i)\;\leq\;0\quad\text{for all }\;i,\hat{t}_i,t_i. \tag{21}$$

Finally, condition (5) follows from Lemma 2. *QED*

Our assumption that any best outcome for the principal is a worst outcome for all agents may possibly be weakened in some cases, because the assumption guarantees that Walras' law holds for every agent's entire demand correspondence in the fictitious economy, whereas in the proof we only use the fact that Walras' law holds at the equilibrium prices. However, the assumption cannot be completely dropped.

Consider *Example 1.* There is only one agent, with no private information, and the principal has two equally likely types, $T_0 = \{0, 1\}$. The space of basic outcomes is the unit interval $Z = [0, 1]$. The players have single-peaked preferences, $u_0(z, t_0) = -(z - t_0)^2$ and $u_1(z) = -z^2$. (Hence, the agent's preferences are aligned with type 0 of the principal.) The disagreement outcome is $z_0 = 1/2$. The following deterministic allocation rule ρ dominates all other incentive feasible allocation rules: $\rho(0) = 0$ and $\rho(1) = \sqrt{2}$. Hence, ρ is a perfect Bayesian equilibrium outcome (assume prior beliefs if any alternative mechanism is proposed, and apply the revelation principle).[18] But ρ is not an SUPO: for the belief Q_0 that puts probability 1/4 on type 1 of the principal, the deterministic allocation rule ρ' given by $\rho'(0) = 0$ and $\rho'(1) = 1$ is Q_0-unconstrained dominating ρ. Hence, an SUPO does not exist.

6 Computing SUPOs in generalized private values environments

Computing slack exchange equilibria can be difficult. However, in many cases a shortcut is possible.

Observe that $\mathcal{Z}$ is a convex subset of the vector space of signed measures on Z. Hence, the allocation rule ρ belongs to the vector space of functions from $\mathbf{T}$ into the space of signed measures on Z. The set of allocation rules is convex and is denoted Ω.

For all $\rho \in \Omega$, let $G(\rho, i, t_i)$ denote the left-hand-side of constraint $(i, t_i) \in \mathcal{R}$ in any problem $P(t_0, r, c)$. Let $G(\rho, i, \hat{t}_i, t_i)$ denote the left-hand-side of constraint $(i, \hat{t}_i, t_i) \in \mathcal{C}$. Let $G(\rho)$ denote the corresponding function $\mathcal{R} \cup \mathcal{C} \rightarrow \mathbb{R}$. Let $f^{t_0}(\rho)$ denote the objective of problem $P(t_0, r, c)$. Let the Lagrange

[18] By Myerson, a neutral optimum exists and is undominated if we replace the outcome space by a finite set Z'; hence, ρ is the unique neutral optimum if $Z' \supseteq \{0, \sqrt{2}\}$.

function be denoted by

$$L^{t_0,r,c}(\rho,\beta,\gamma) \quad = \quad f^{t_0}(\rho) + (\beta,\gamma)\cdot(G(\rho)+(r,c)).$$

The *saddle-point condition* is satisfied at $(\hat{\rho},\hat{\beta},\hat{\gamma})$ if, for all $\rho \in \Omega$ and all non-negative (β,γ),

$$L^{t_0,r,c}(\rho,\hat{\beta},\hat{\gamma}) \leq L^{t_0,r,c}(\hat{\rho},\hat{\beta},\hat{\gamma}) \leq L^{t_0,r,c}(\hat{\rho},\beta,\gamma).$$

The vector $(\hat{\beta},\hat{\gamma})$ is then called a Lagrange multiplier vector.

We say that the vectors (β,γ) and (β',γ') are *co-linear* if there exists $k>0$ such that the $\beta(i,t_i) = k\beta'(i,t_i)$ for all $(i,t_i) \in \mathcal{R}$ and $\gamma(i,\hat{t}_i,t_i) = k\gamma'(i,\hat{t}_i,t_i)$ for all $(i,\hat{t}_i,t_i) \in \mathcal{C}$.

Lemma 3. *Consider a generalized private values environment. Suppose that, for all t_0 and some $(r^*_{t_0},c^*_{t_0})$ and $(\hat{\beta},\hat{\gamma})$, the saddle-point condition is satisfied for problem $P(t_0,r^*_{t_0},c^*_{t_0})$ at the point $(\rho,\hat{\beta},\hat{\gamma})$. Moreover, suppose that (5) holds with $(\beta^*,\gamma^*) = (\hat{\beta},\hat{\gamma})$.*

Then (4) holds for all (β^*,γ^*) that are co-linear to $(\hat{\beta},\hat{\gamma})$. Moreover, ρ is a maximizer of problem $P(t_0,r^*_{t_0},c^*_{t_0})$.

Proof. From Luenberger ((1969), p. 221, Theorem 2), ρ is a maximizer of problem $P(t_0,r^*_{t_0},c^*_{t_0})$. From Luenberger ((1969), p. 222, Theorem 1),

$$V(t_0,r^*_{t_0},c^*_{t_0}) - V(t_0,r,c) \quad \geq \quad (\hat{\beta},\hat{\gamma})\cdot\left((r^*_{t_0},c^*_{t_0})-(r,c)\right).$$

From (5) with $(\beta^*,\gamma^*) = (\hat{\beta},\hat{\gamma})$, we have $(\hat{\beta},\hat{\gamma})\cdot(r^*_{t_0},c^*_{t_0}) = 0$. Given the constraint $(\hat{\beta},\hat{\gamma})\cdot(r,c) \leq 0$ of the problem in (4) with $(\beta^*,\gamma^*) = (\hat{\beta},\hat{\gamma})$, we conclude that $V(t_0,r^*_{t_0},c^*_{t_0}) \geq V(t_0,r,c)$. Hence, $(r^*_{t_0},c^*_{t_0})$ solves the problem in (4) with $(\beta^*,\gamma^*) = (\hat{\beta},\hat{\gamma})$, and hence with any co-linear (β^*,γ^*). *QED*

This lemma is the key towards computing slack exchange equilibria. One determines slacks $(r^*_{t_0},c^*_{t_0})$ for all t_0 such that the Lagrange multiplier vectors $(\beta_{t_0},\gamma_{t_0})$ for the problems $P(t_0,r^*_{t_0},c^*_{t_0})$ become co-linear across types t_0, and such that the budget conditions (6) and (7) are satisfied.

In best-worst environments condition (5) is then automatically satisfied and we have found a slack exchange equilibrium.

Lemma 3 shows that the actual process of solving the optimization problem (4) can be avoided if the saddle-point condition is satisfied. This shows that the technique of working with co-linear Lagrange multipliers generalizes tremendously beyond the environments considered by Maskin and Tirole (1990).

We obtain particularly simple necessary and sufficient conditions for no trade being an equilibrium,$(r^*_{t_0}, c^*_{t_0}) \leq 0$ for all t_0, because here the budget constraints (6) and (7) are automatically satisfied.

Proposition 4. *Consider any generalized private values constraint-non-degenerate environment, where any best outcome for the principal is a worst outcome for all agents. Furthermore, let $\hat{\rho}$ be an allocation rule that would be optimal for the principal if her type were known to the agents.*

Then, $\hat{\rho}$ is a slack exchange equilibrium if and only if there exists a set of co-linear $(\beta^*_{t_0}, \gamma^*_{t_0})$ such that the saddle-point condition is satisfied for problem $P(t_0, 0, 0)$ at the point $(\hat{\rho}, \beta^*_{t_0}, \gamma^*_{t_0})$ for all $t_0 \in T_0$.

The proof is given in the supplement to this paper.

We now compute two examples of slack exchange equilibria. Our first example has private values with linear risk-neutral payoff functions and is a more general version of the example considered in Section 2. This example can be interpreted as a discrete types version of partnership dissolution problem in Cramton, Gibbons, and Klemperer (1987) or as a version of the speculative trade environment with non-common priors considered in Eliaz and Spiegler (2006, 2007).

In this example, in equilibrium only *one* constraint is traded at 0 price. Furthermore, the aggregate consumption of the slack is negative for this constraint and 0 for *all* other constraints. Finally, the principal's type that buys the slack achieves a higher payoff than would be possible for her if the types of the players were commonly known.

There are a principal and an agent. The set of outcomes is given by

$$Z = [-1, 1] \times [\underline{x}, \overline{x}].$$

We assume that $T_0 = \{0, 3\}$ and $T_1 = \{1, 2\}$ and that $\underline{x} \leq -2$ and $\overline{x} \geq 2$. The prior probabilities of the types are common knowledge; the prior probability of $t_j^{(i)}$, where j is the player's index and i is the type's index, is denoted by $p_j^{(i)}$. We require that $p_j^{(i)} > 0$ for any j and i, $p_0^{(1)} \geq p_0^{(2)}$, and $p_1^{(1)} \geq p_1^{(2)}$. Let $\alpha = p_0^{(2)}/p_0^{(1)}$. The disagreement outcome is $z_0 = (0, 0)$.

The players have linear risk-neutral payoff functions with private values:

$$\begin{aligned} u_0(y, x, t_0) &= x - yt_0, \\ u_1(y, x, t_1) &= yt_1 - x. \end{aligned}$$

We now describe a slack exchange equilibrium:

$$\rho^* = (y^*, x^*)(t_0, t_1) \quad = \quad \begin{cases} (1, 1), & \text{if } t_0 = 0; \\ (-1, -1), & \text{otherwise.} \end{cases}$$

Let $(r^*_{t_0}, c^*_{t_0})$ denote the amount of slack corresponding to ρ^* and observe that

$$r^*_{t_0}(1,1) = 0,\ r^*_{t_0^{(1)}}(1,2) = -r^*_{t_0^{(2)}}(1,2) = 1,\ c^*_{t_0} = 0,$$

Furthermore, define

$$\beta^*_1 = 1, \quad \beta^*_2 = 0, \quad \gamma^*_{21} = p_1^{(2)}, \quad \gamma^*_{12} = 0.$$

Remark 1. *The allocation rule ρ^* is a maximizer of program $P(t_0, r^*_{t_0}, c^*_{t_0})$. Furthermore, $(r^*_{t_0}, c^*_{t_0})_{t_0 \in T_0}, \beta^*, \gamma^*$ is the slack exchange equilibrium corresponding to ρ^*.*

The proof is given in the supplement to this paper.

Our second example is an environment with generalized private values and linear risk-neutral payoff functions. The special cases of this environment include a discrete-type version of the Myerson-Satterthwaite bargaining environment where the parties have private information about their valuations of the good, a version of the Akerlof's Lemons market where the seller (agent) is privately informed about the quality of the good and the buyer (principal) is privately informed about its willingness to pay for additional quality, and a labor contract setting in which the worker (agent) has private information about its productivity and the firm (principal) has private information about demand for its product.

In this environment, the set of constraints that cannot be relaxed and, hence, the *set* of Lagrange multiplier vectors for which the saddle point condition is satisfied varies with the type of the principal. Nevertheless, it is possible to find *a* Lagrange multiplier vector such that the saddle-point condition is satisfied for all principal types. Thus, the solution of the informed principal problem coincides with the allocation rules that would be implemented in the absence of uncertainty about the principal's type.

There are a principal and an agent. The set of outcomes is given by

$$Z \;=\; [0,1]\times[\underline{x},\overline{x}],$$

where any $y \in [0,1]$ represents the amount of good allocated to the agent and any $x \in [\underline{x},\overline{x}]$ represents a monetary transfer from the agent to the principal.

We assume that the agent's type space is finite, $T_1 = \{1,\ldots,k\}$. Furthermore, we require that $k < \overline{x}, -\underline{x}$ and $1 \le t_0 \le k$ for all $t_0 \in T_0$. The prior probability of $t_1 = i$ is denoted by $p_1^{(i)}$. Let $z_0 = (0,0)$ be the disagreement outcome.

The players have linear risk-neutral payoff functions with generalized private values:

$$\begin{aligned} u_0(y,p,t_0,t_1) &= (1-y)f(t_0,t_1)+x,\\ u_1(y,p,t_1) &= yt_1 - x, \end{aligned}$$

where $f(t_0,t_1) > 0$ for all $t_0 \in T_0$ and $t_1 \in T_1$.

Define the *virtual surplus function*

$$v(i) \;=\; \begin{cases} k, & \text{if } i = k;\\ i - \frac{G^{(i)}}{p_1^{(i)}}, & \text{otherwise;} \end{cases}$$

where $G^{(i)} = 1 - \sum_{j=1}^{i} p_1^{(i)}$ for all $i = 1,\ldots,k-1$.

In this environment, any allocation rule ρ can be decomposed into a *good allocation rule* $\mu : \mathbf{T} \to [0,1]$ and a *transfer allocation rule* $\tau : \mathbf{T} \to [\underline{x},\overline{x}]$; that is, $\rho = (\mu,\tau)$. Let

$$\begin{aligned} \mu^*(t_0,i) &= \begin{cases} 1, & \text{if } v(i) - f(t_0,i) \ge 0;\\ 0, & \text{otherwise;} \end{cases}\\ \tau^*(t_0,i) &= \begin{cases} \sum\limits_{j=2}^{i}(\mu^*(t_0,j)-\mu^*(t_0,j-1))\cdot j + \mu^*(t_0,1), & \text{if } i > 1;\\ \mu^*(t_0,1), & \text{if } i = 1. \end{cases} \end{aligned}$$

Proposition 5. *Let $v(i) - f(t_0,i)$ be increasing in i for all t_0. Then, the allocation rule $\rho^* = (\mu^*,\tau^*)$ is a maximizer of the problem $P(t_0,0,0)$ for any t_0 and is a slack exchange equilibrium.*

The proof is given in the supplement to this paper.

References

Akerlof, G. A. (1970): "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism," *The Quarterly Journal of Economics*, 84(3), 488–500.

Cramton, P., R. Gibbons, and P. Klemperer (1987): "Dissolving a Partnership Efficiently," *Econometrica*, 55(3), 615–32.

Debreu, G. (1959): *Theory of Value: An axiomatic analysis of economic equilibrium*. Yale University Press, New Haven and London.

Eliaz, K., and R. Spiegler (2006): "Contracting with Diversely Naive Agents," *Review of Economic Studies*, 73(3), 689–714.

——— (2007): "A Mechanism-Design Approach to Speculative Trade," *Econometrica*, 75(3), 875–884.

Guesnerie, R., and J.-J. Laffont (1984): "A complete solution to a class of principal-agent problems with an application to the control of a self-managed firm," *Journal of Public Economics*, 25(3), 329–369.

Luenberger, D. G. (1969): *Optimization by Vector Space Methods*. John Wiley and Sons, Inc., New York, NY.

Maskin, E., and J. Tirole (1986): "Principals with private information, I: Independent values," Harvard University Discussion Paper # 1234.

——— (1990): "The principal-agent relationship with an informed principal: The case of private values," *Econometrica*, 58(2), 379–409.

——— (1992): "The principal-agent relationship with an informed principal, II: Common values," *Econometrica*, 60(1), 1–42.

Myerson, R. B. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6(1), 58–73.

——— (1983): “Mechanism design by an informed principal,” *Econometrica*, 51(6), 1767–1798.

Myerson, R. B., and M. A. Satterthwaite (1983): “Efficient mechanisms for bilateral trading,” *Journal of economic theory*, 29(2), 265–281.

Mylovanov, T., and T. Tröger (2008): “Optimal Auction Design and Irrelevance of Privacy of Information,” *mimeo*.

Yilankaya, O. (1999): “A note on the seller’s optimal mechanism in bilateral trade with two-sided incomplete information,” *Journal of Economic Theory*, 87(1), 125–143.