Games and Economic Behavior 94 (2015) 169–181

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Optimal private good allocation: The case for a balanced budget

Moritz Drexl, Andreas Kleiner [*]

Bonn Graduate School of Economics, University of Bonn, Kaiserstr. 1, 53113 Bonn, Germany

ARTICLE INFO

Article history:
Received 11 December 2012
Available online 28 October 2015

JEL classification:
D82
D45

Keywords:
Mechanism design
Bilateral trade
Myerson–Satterthwaite theorem
Budget balance

ABSTRACT

In an independent private value auction environment, we are interested in strategy-proof mechanisms that maximize the agents' residual surplus, that is, the utility derived from the physical allocation minus transfers accruing to an external entity. We find that, under the assumption of an increasing hazard rate of type distributions, an optimal deterministic mechanism never extracts any net payments from the agents, that is, it will be budget-balanced. Specifically, optimal mechanisms have a simple "posted price" or "option" form. In the bilateral trade environment, we obtain optimality of posted price mechanisms without any assumption on type distributions.

1. Introduction

Most parts of the mechanism design literature studying welfare maximization problems focus on mechanisms implementing the efficient allocation. However, in general it is not possible to implement the efficient allocation in dominant strategies using budget-balanced mechanisms (Green and Laffont, 1979). Given this result, we study how to choose among different mechanisms that cannot attain both, allocative efficiency and budget-balancedness. Since we are concerned with welfare maximization, the social planner's objective function should consist of the agents' aggregate utility and therefore include aggregate transfers. In other words, one seeks to find mechanisms that maximize what we call residual surplus. This is the surplus, or utility, the agents derive from the chosen physical allocation, reduced by the amount of transfers that are lost to an external agency (this is often called "money burning").

A common approach is to implement the efficient allocation via Groves mechanisms and to redistribute as much money to the agents as possible without distorting incentives (Cavallo, 2006; Guo and Conitzer, 2009, 2010; Moulin, 2009). This approach aims at characterizing the optimal mechanism for allocating private goods that implements the *efficient allocation* in *dominant strategies*, is *individually rational* and *never creates a budget deficit* (ex-post).[1] However, if mechanisms that allocate inefficiently yield higher residual surplus (Guo and Conitzer, 2014) it is not clear why one should use a mechanism that allocates efficiently.

Consequently, we drop the requirement that mechanisms allocate efficiently. Instead, we take an optimal mechanism design approach and consider mechanisms that are comparable to the ones considered before in that they are *strategy-proof*, *deterministic*, *never run a deficit* and satisfy *ex-post participation constraints*. We analyze which mechanism maximizes residual

* Corresponding author.
E-mail addresses: drexl@uni-bonn.de (M. Drexl), akleiner@uni-bonn.de (A. Kleiner).

[1] In our setting, the best Groves mechanism is implemented by a second-price auction with two bidders.

http://dx.doi.org/10.1016/j.geb.2015.10.009
0899-8256/

surplus when an indivisible good is auctioned among two agents with independent private values that are distributed according to prior type distributions. We show that under an increasing hazard rate assumption on type distributions, the optimal mechanism will never waste any payments, thereby deviating distinctly from the efficient allocation (Theorem 1). In fact, our proof method reveals that all mechanisms that allocate efficiently are worse than the simple mechanism where the object is always given to the same agent (the one with the higher expected valuation; Corollary 1), showing that our general mechanism design approach has clear advantages over the previous approach to search for the optimal Groves mechanism. We show that the optimal mechanism is either a "posted price" or an "option" mechanism: The object is assigned to one of the agents unless both agents agree to trade at a prespecified price (posted price mechanism) or unless the second agent uses his option to buy the object at a fixed price from the first agent (option mechanism). Therefore, the optimal mechanisms do not invoke money burning and are of a particularly simple form. Moreover, numerical simulations indicate that these simple mechanisms obtain a large share of first-best welfare (92 per cent on average in our simulations). In the bilateral trade setting, we establish optimality of posted price mechanisms without any restrictions on type distributions (Theorem 2). This provides an argument for the focus on budget-balanced mechanisms (see Myerson and Satterthwaite, 1983; Hagerty and Rogerson, 1987).

The requirement that a mechanism does not produce a budget deficit ex-post is considerably stronger than the requirement that this holds in expectation. However, in many situations it is reasonable that a budget breaker is infeasible and therefore ex-post constraints need to be obeyed. This includes situations where there is no insurance or where agents have restricted access to capital markets. Also, hidden information issues towards a third party cannot always be resolved, and autarkic mechanisms that can be implemented without explicit intervention by a third party might be preferable (e.g., when mechanisms are used to model bargaining situations; Myerson and Satterthwaite, 1983; Hagerty and Rogerson, 1987). If all these considerations do not apply and mechanisms that create no deficit in expectation can be implemented (for example, because the designer has unlimited liability), then one can achieve the first-best solution (see Section 5). Similarly, we show that one can achieve the first-best if mechanisms are only required to be Bayesian incentive compatible (Proposition 1). In contrast to these two constraints, which are the main driving forces behind our results, we argue that the participation constraint and the restriction to deterministic mechanisms are not essential to the spirit of our results (Section 5).

Our work is part of a small literature that searches for mechanisms maximizing residual surplus when the first-best is not achievable. Miller (2012) studies a model of firms colluding in a Bertrand oligopoly. A mechanism used by a cartel to allocate market shares should maximize residual surplus. Miller shows that under general conditions it is never optimal to allocate market shares efficiently and gives numerical evidence that for some type distributions it is optimal to give up efficiency in order to obtain a balanced budget. However, other examples indicate that this observation does not hold for all distributions. Athey and Miller (2007) study residual surplus maximization in a repeated bilateral trade setting and obtain numerical results suggesting that for many type distributions the optimal mechanism is a posted price mechanism. Closely related to our paper is independent work by Shao and Zhou (2012), who obtain the characterization of our Theorem 1 when restricting to symmetric distributions of types and allowing mechanisms to violate individual rationality.

The result that the efficient allocation is never optimal contrasts with the literature cited above that restricts attention to efficient rules (Cavallo, 2006; Guo and Conitzer, 2009, 2010; Moulin, 2009). Recently, Athanasiou (2013) and Sprumont (2013) relax this requirement. Similarly to our work, they focus on mechanisms that are deterministic, strategy-proof, ex-post individually rational and create no deficit ex-post. However, they require mechanisms in addition to be anonymous, which immediately implies that whenever the object is allocated, it is allocated to the agent that values it the most (weak assignment efficiency). This restricts the set of mechanisms severely and excludes the mechanisms that turn out to be optimal in our analysis.

The restriction to efficient allocation rules has also been relaxed in a series of papers that study specific mechanisms in a multi-unit setting. Faltings (2005) and Moulin (2009) propose simple mechanisms where one agent is designated as a residual claimant and is allocated one unit (or no unit, respectively) independent of his type. The remaining units are auctioned among the other agents and the residual claimant receives all payments accruing in the auction. Faltings uses numerical examples to argue that his mechanism often outperforms the VCG mechanism. Moreover, Moulin shows that his mechanism provides a higher worst-case welfare guarantee than any VCG mechanism given that there are sufficiently many objects and agents. In our setting with two agents and one object, these mechanisms always allocate the object to a fixed agent and therefore correspond to a degenerate option mechanism. Our Corollary 1 supports Faltings' numerical results in the two agent setting by showing that under regular prior distributions his mechanism indeed outperforms the VCG mechanism. Building on the ideas of Faltings and Moulin, Guo and Conitzer (2014) provide worst-case welfare guarantees for two specific classes of mechanisms that allocate inefficiently: *Burning allocation mechanisms* burn a (random) number of units and assign the remaining units efficiently. *Partitioning mechanisms* partition units and agents randomly into two groups, allocate the objects in each partition efficiently to the agents in the corresponding partition and distribute the payments to agents in the other partition. Similarly, de Clippel et al. (2014) propose a deterministic mechanism where the burning of items is contingent on the reports of the agents; they provide worst-case welfare guarantees that converge to 0.88 asymptotically as the number of agents grows. Our work differs from these papers by evaluating mechanisms according to a Bayesian prior, restricting ourselves to the two agent setting and using a general optimal mechanism design perspective.

Another related strand of the literature studies the expected residual surplus of Bayesian incentive compatible mechanisms when it is not possible to redistribute any payments among the agents (Hartline and Roughgarden, 2008; Chakravarty and Kaplan, 2013; Condorelli, 2012). This implies that methods similar to those in Myerson (1981) can be applied. It is shown that for a large class of type distributions (those which exhibit an increasing hazard rate) it is optimal to always assign the object to the same agent. Maximization of residual surplus also plays a role in the analysis of optimal mechanisms used by bidding rings (McAfee and McMillan, 1992). It is worth noting that the equivalence between Bayes–Nash and dominant strategy implementation (Gershkov et al., 2013; Manelli and Vincent, 2010) does not apply to our model.[2]

We present the basic model for the auction environment in Section 2 and characterize incentive compatible mechanisms in Section 3. The optimization problem is solved in Section 4, the role of the assumptions is discussed in Section 5. We study this mechanism design problem in the bilateral trade context in Section 6, and conclude in Section 7.

2. Model

An indivisible object is auctioned among two agents. Each agent $i = 1, 2$ has a valuation x_i for the object, which is his private information. Valuations are drawn independently from $X_i = [0, \bar{x}_i]$ according to distribution functions F_i with corresponding densities f_i, which we assume to be bounded.[3] We denote by $X = X_1 \times X_2$ the product type space and by F the joint distribution on X. For notational convenience, when concentrating on agent i, we will write (x_i, x_{-i}) for $x = (x_1, x_2) \in X$.

If agent i is given a payment of p_i (usually negative), his utility is $x_i + p_i$ for winning the object, and p_i if the other agent gets the object.

Mechanisms Due to the Revelation Principle we focus on truthfully implementable direct revelation mechanisms for selling the object.

Definition 1. A *mechanism* M is a tuple (d, p), where $d : X \to \{0, 1\}^2$ and $p : X \to \mathbb{R}^2$ are measurable functions, such that $d_1(x) + d_2(x) = 1$.[4]

The interpretation is that $d_i(x) = 1$ if and only if agent i gets the object. If the agents report x, then agent i receives as payment the component $p_i(x)$ of $p(x)$.

Equilibrium concept We consider strategy-proof mechanisms where truthful reporting is a dominant strategy for both agents. Thereby, we ensure that the mechanisms can robustly be implemented without specific assumptions on the beliefs of the agents. Hence, we define the following notion of incentive compatibility:

Definition 2. A mechanism M is *incentive compatible (IC)* if for every agent i and for each $x_i \in X_i$, $r_i \in X_i$,

$$d_i(x_i, r_{-i}) \cdot x_i + p_i(x_i, r_{-i}) \geq d_i(r_i, r_{-i}) \cdot x_i + p_i(r_i, r_{-i})$$

holds for each $r_{-i} \in X_{-i}$.

This definition is independent of the distribution of valuations, which reflects the robustness of strategy-proof mechanisms as compared to mechanisms that are Bayes–Nash incentive compatible. Although the set of mechanisms we consider does therefore not depend on F, the next section shows that the distributions determine which mechanism is optimal.

Objective and further constraints We aim at finding the mechanism that maximizes the sum of agents' ex-ante (expected) residual surplus, that is, utility derived from the physical allocation minus aggregate payments. We impose the constraint that the mechanism has to be *ex-post no-deficit (ND)*, that is, for every type profile x, we require $p_1(x) + p_2(x) \leq 0$.[5] Also, the mechanism has to be *ex-post individually rational (IR)*, that is, for all type profiles x, we require $d_i(x)x_i + p_i(x) \geq 0, i = 1, 2$. Summarizing, we want to solve the following optimization problem:

$$\max_{M=(d,p)} \int_X \Big[d_1(x)x_1 + d_2(x)x_2 + p_1(x) + p_2(x)\Big] dF(x)$$
$$\text{s.t. } M \text{ satisfies IC, ND and IR.} \tag{1}$$

We say that a mechanism is optimal if it solves problem (1).

[2] See Section 5 for more details.

[3] Assuming that the lower bound of the type space is 0 simplifies the analysis. The details are explained in footnote 7.

[4] For a discussion of stochastic mechanisms, see Section 5. We follow Athey and Miller (2007) and Miller (2012) and assume that the good is always allocated. This is reasonable, for example, when considering how a cartel allocates market shares, or how the government sells licenses to firms. While there can be welfare gains from not allocating the good when one focuses on anonymous mechanisms (de Clippel et al., 2014), these gains seem to be minor in our model. Moreover, the assumption that the good is always allocated is without loss of generality in the trade setting (Section 6).

[5] Ex-post budget constraints are commonly imposed on mechanism design problems: see, for example, the literature on optimal redistribution (Guo and Conitzer, 2010, 2009; Moulin, 2009) and bilateral trade (Hagerty and Rogerson, 1987; Myerson and Satterthwaite, 1983), or Chawla et al. (2006). The role of this assumption is discussed in Section 5.

3. Characterization of incentive compatibility

The aim of this section is to give a characterization of incentive compatibility in order to simplify problem (1). The conditions characterizing incentive compatible mechanisms involve a monotonicity and an integrability condition. We first define monotonicity.

Definition 3. The allocation function d is *monotone* if d_i is non-decreasing in x_i for $i = 1, 2$.

Now given a monotone allocation function d, define the following functions for $i = 1, 2$:

$$g_i(x_{-i}) := \inf\{x_i : d_i(x_i, x_{-i}) = 1\}.$$

If there is no x_i such that $d(x_i, x_{-i}) = 1$, then we set $g_i(x_{-i}) = \bar{x}_i$. Note that if d is monotone, these functions define d almost everywhere. The following lemma, which is a corollary of Myerson (1981), gives a characterization of incentive compatibility.

Lemma 1. *A deterministic mechanism $M = (d, p)$ is incentive compatible if and only if the following two conditions are satisfied:*

1. *The allocation rule d is a monotone step function.*
2. *For all $x_i \in X_i$ and $x_{-i} \in X_{-i}$,*

$$p_i(x_i, x_{-i}) = q_i(x_{-i}) - g_i(x_{-i}) d_i(x_i, x_{-i}) \tag{2}$$

for some function $q_i : X_{-i} \to \mathbb{R}$.

The interpretation of condition (2) is that an agent who receives the object is punished by receiving a lower payment: she receives $q_i(x_{-i})$ if she does not receive the object, and this payment is reduced in case she gets the object to make the agent's marginal type $g_i(x_{-i})$ indifferent between receiving and not receiving the object.

This can be interpreted as a payoff-equivalence result: Payments are completely determined by the allocation as soon as one fixes the payment for some type x_i. Hence, the only freedom that is left regarding the payment scheme, is to give the agent an additional payment that is independent of his type. These additional payments can serve as a possibility to redistribute certain amounts of payments to another agent. Given an allocation rule d and a payment rule p, we say that the redistribution payment q implicitly defined by the above equality is associated with p.

The simplified formulation of problem (1) is the following:

$$\max_{M=(d,p)} \int_X \Big[d_1(x)[x_1 - g_1(x_2)] + d_2(x)[x_2 - g_2(x_1)] + q_1(x_2) + q_2(x_1) \Big] dF(x)$$

$$\text{s.t. } M \text{ satisfies IR and ND, } q \text{ is associated with } p \text{ and } d \text{ is monotone.}$$

We will write $U(M)$ for the above integral and from now on only consider mechanisms that are IC, IR and ND.

4. The optimal auction

In this section, we present the first main result of this paper: if we impose an increasing hazard rate condition on the type distributions, then the optimal mechanism is always budget-balanced. Specifically, it turns out that the optimal mechanism takes one of two simple forms:

Either it is a posted price mechanism which by default allocates the object to one of the agents (agent 1, say) and changes the allocation if and only if both agents agree to trade at a prespecified price a, i.e., agent 1 reports a valuation below a fixed price a and agent 2 reports a valuation above a. If agent 2 is allocated the object, he makes a payment a to agent 1, otherwise no transfers accrue.

Or it is an option mechanism where the good is allocated by default to agent 1, but agent 2 has the option to buy the object at price a. Hence, if agent 2's valuation is above the strike price a, he buys the object and pays a to agent 1 (see also Shao and Zhou, 2012).

Formally, these two mechanisms are defined as follows:

Definition 4. A mechanism $M = (d, p)$ is a *posted price mechanism with default agent 1 and price a*, if

$$\begin{array}{ll} d_2(x) = 1,\ p(x) = (a, -a) & \text{if } x_1 \le a \text{ and } x_2 \ge a, \\ d_2(x) = 0,\ p(x) = (0, \quad 0) & \text{otherwise.} \end{array}$$

M is an *option mechanism with default agent 1 and price a*, if

$$\begin{array}{ll} d_2(x) = 1,\ p(x) = (a, -a) & \text{if } x_2 \ge a, \\ d_2(x) = 0,\ p(x) = (0, \quad 0) & \text{otherwise.} \end{array}$$

Similarly, one can define posted price and option mechanisms with default agent 2. If we do not specify the agent or price we just say that M is option or posted price.

Both classes of mechanisms are parameterized by the price a and it is easy to check that all these mechanisms are budget-balanced as well as incentive compatible and individually rational.

Our assumption on type distributions is the following:

Condition *(HR). The hazard rates of the type distributions are monotone. That is, the functions* $h_i(x_i) = \frac{f_i(x_i)}{1-F_i(x_i)}$ *are non-decreasing in* $x_i \in [0, \bar{x}_i)$ *for* $i = 1, 2$.

Theorem 1. *Suppose that the hazard rates of the type distributions are monotone. Then the optimal mechanism is either a posted price or an option mechanism.*

It is known that if payments are wasteful by assumption, then for regular distributions it is optimal to make the allocation independent of reports (Hartline and Roughgarden, 2008): a more efficient allocation is more than offset by the waste of payments that are required for incentive-compatibility. Given that money can be redistributed in our model, there are better budget-balanced mechanisms (essentially posted price and option mechanisms). One might argue naively that, if wasting money is suboptimal in the setting of (Hartline and Roughgarden, 2008), it must also be suboptimal in this setting, and hence a budget-balanced mechanism must be optimal. However, redistribution payments allow for additional flexibility, which makes the argument more subtle and requires that we optimize jointly allocations and redistribution payments.[6]

The proof can be sketched as follows: We first show the important auxiliary result that either an option mechanism or a posted price mechanism is optimal in $\mathcal{M}_0$, the class of mechanisms such that g_i is monotone and piecewise constant for each agent (Lemma 2). We then argue that the residual surplus $U(M)$ of a given mechanism M can be approximated arbitrarily well by a mechanism in $\mathcal{M}_0$ (Lemma 3). The Theorem then follows by the following observation: Suppose there is a mechanism $\bar{M}$ being strictly better than the best option or posted price mechanism, and denote the difference in residual surplus by ε. It follows from Lemma 3 that there is a mechanism in the class $\mathcal{M}_0$ whose residual surplus is within $\frac{\varepsilon}{2}$ of $U(\bar{M})$, thus being better than the best option or posted price mechanism. But this contradicts Lemma 2, hence there cannot be a mechanism being better than the best option or posted price mechanism.

While the approximation part of the proof can be found in the Appendix, we state and prove Lemma 2, which contains the essence of why Theorem 1 holds.

Lemma 2. *Suppose that the hazard rates of the type distributions are monotone and let* $M = (d, p)$ *be any mechanism in* $\mathcal{M}_0$*. Then there exists a mechanism* M' *that is posted price or option such that* $U(M') \geq U(M)$.

Proof. The proof consists of three steps: *Step 1* determines for an arbitrary allocation rule the maximal possible redistribution payments q_i. Hence, the allocation rule from this point on completely determines the optimal payments and we can constructively manipulate the allocation rule in *Steps 2* and *3* until we end up with an option or posted price mechanism.

Step 1: We denote the jump points of $g_2(x_1)$ and $g_1(x_2)$ by α_j and β_j, respectively (see Fig. 1). Note that, without loss of generality, we can assume that for the first segment of g_1 we have $g_1(x_2) = 0$ since otherwise we could switch the roles of the agents.

We now claim that $q_2(x_1) = 0,\ \forall x_1 \in X_1$; that is, no money is redistributed to agent 2. To see this, pick arbitrary x_1 and observe that $g_1(0) = 0$[7] and $d_2(x_1, 0) = 0$; therefore $g_1(0)d_1(x_1, 0) = g_2(x_1)d_2(x_1, 0) = 0$. From (ND) it follows that $q_1(0) + q_2(x_1) = p_1(x_1, 0) + p_2(x_1, 0) \leq 0$. Also, (IR) for agent 2 at $(x_1, 0)$ implies $q_2(x_1) \geq 0$, and (IR) for agent 1 at $(0, 0)$ implies $q_1(0) \geq 0$, and therefore $q_2(x_1) = 0$.

Next, we can assume that

$$q_1(x_2) = \min_{x_1} \left\{ g_1(x_2)d_1(x_1, x_2) + g_2(x_1)d_2(x_1, x_2) \right\} \tag{3}$$

always holds, since by (ND) this relation always holds with $\leq$ and changing it to equality does not reduce $U(M)$. In this way, the complete payment-scheme is determined through the allocation rule d. Note that setting the function q this way implies that (ND) and (IR) are always satisfied.

Step 2: In this step we argue that changing the allocation to the one shown in Fig. 1b does not increase money burning, but increases allocative efficiency and hence aggregate welfare.

[6] Indeed, if the arguments from a model without redistribution could simply be extended, our conclusion would also hold for stochastic mechanisms. However, numerical results in Section 5 show that this is not the case.

[7] At this step we use that 0 is the lower bound of the type space. This assumption implies that participation constraints pin down the maximal redistribution payments. Without this assumption, one would have to optimize over redistribution payments. If, for example, $f_1 = f_2$ and f_i is log-concave then exactly the same results can be obtained.

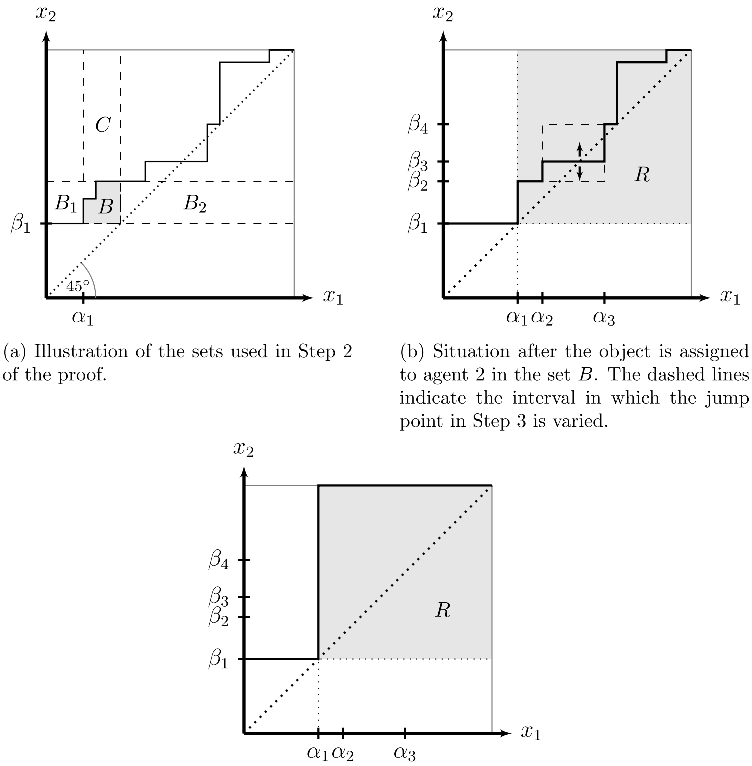

(a) Illustration of the sets used in Step 2 of the proof.

(b) Situation after the object is assigned to agent 2 in the set B. The dashed lines indicate the interval in which the jump point in Step 3 is varied.

(c) A posted price mechanism with default agent 1 and price β_1 improves over M.

Fig. 1. Illustration of the proof of Lemma 2.

Define the set $B = \{x \mid x_1 \leq \beta_1 \leq x_2, d_2(x) = 0\}$ and consider the sets B_1, B_2 and C as shown in Fig. 1a. We change the allocation rule and allocate the object to agent 2 for types in B. Since $x_2 \geq x_1$ for $x \in B$, this improves the physical allocation and we can concentrate on payments. Note that q_1, as defined in (3), increases to the same extent as g_1, hence any additional payments in the set B_2 can be redistributed. Also, transfers are weakly increased for types in B_1 and C. As the change in allocation has no effect outside these sets, the claim follows.

Step 3: This step studies the effects of shifting steps in the set R, shown as the shaded area in Fig. 1b, while fixing redistribution payments. Our condition on the hazard rate ensures that each step should optimally be moved to either the lowest or the highest possible position. Hence, proceeding iteratively, we obtain either an option mechanism or a posted price mechanism. This will complete the proof.

Changing the allocation in R does not change q_1 as defined in (3) and we ignore the functions q_i from now on.

The following is a procedure to remove one step contained in R without decreasing $U(M)$. We do this exemplarily with the jump point at β_3 (see Fig. 1b). We vary β_3 on the interval $[\beta_2, \beta_4]$ and show that welfare is quasi-convex in β_3. This implies that setting $\beta_3^* = \beta_2$ or β_4 increases $U(M)$. The part of $U(M)$ that depends on β_3 is the following:

$$\int_{\alpha_2}^{\alpha_3} \left[\int_{\beta_2}^{\beta_3} (x_1 - \alpha_2) dF_2(x_2) + \int_{\beta_3}^{\bar{x}_2} (x_2 - \beta_3) dF_2(x_2) \right] dF_1(x_1) - \int_{\alpha_3}^{\bar{x}_1} \left[\int_{\beta_2}^{\beta_3} \alpha_2 dF_1(x_1) + \int_{\beta_3}^{\beta_4} \alpha_3 dF_2(x_2) \right] dF_1(x_1)$$

Differentiating with respect to β_3 using Leibniz' rule yields

$$\int_{\alpha_2}^{\alpha_3} \Big[f_2(\beta_3)(x_1 - \alpha_2) - \big[1 - F_2(\beta_3)\big] \Big] dF_1(x_1) + \int_{\alpha_3}^{\bar{x}_1} f_2(\beta_3)[\alpha_3 - \alpha_2] dF_1(x_1).$$

Writing constants C_1, C_2 and C_3 for the terms that do not depend on β_3, we get

$$C_1 f_2(\beta_3) - C_2\big[1 - F_2(\beta_3)\big] + C_3 f_2(\beta_3).$$

Table 1
Simulation results showing the share of first-best welfare that is obtained by the optimal posted price or option mechanism. We solved a discretized model and ran 500 trials for each distribution using randomly drawn parameters.

Distribution	Average share	Minimum share
Weibull (IHR)	0.930%	0.788%
Weibull (DHR)	0.916%	0.753%
Exponential	0.905%	0.752%

Assuming $C_2[1 - F_2(\beta_3)] > 0$ (if either $C_2 = 0$ or $1 - F_2(\beta_3) = 0$, we set $\beta_3^* = \beta_4$ without reducing U), we can divide by $C_2[1 - F_2(\beta_3)]$ and get that the derivative is non-negative if and only if

$$C \cdot h_2(\beta_3) - 1 \geq 0,$$

where $C = (C_1 + C_3)/C_2 > 0$. Because $h_2(\beta_3)$ is non-decreasing by condition (HR), quasi-convexity follows and $U(M)$ is increased by either setting $\beta_3^* = \beta_2$ or $\beta_3^* = \beta_4$. In either case, we have decreased the number of steps by one and the procedure ends.

Iteratively applying this procedure establishes the lemma. □

A consequence of the theorem is that, given the increasing hazard rates of the agents' type distributions, finding the best mechanism reduces to finding the best posted price and option mechanisms and comparing these two. For example, if the agents have the same distribution function, all option and posted price mechanisms with the same strike price yield the same welfare and therefore the best mechanism is characterized by the strike price a^* satisfying

$$a^* = \mathbb{E}[x_1] = \mathbb{E}[x_2].$$

Our intermediate results (see the proof of Lemma 2) also allow for a refined judgment of the welfare implied by the efficient allocation, which is employed by the literature on optimal redistribution (Cavallo, 2006; Guo and Conitzer, 2009, 2010; Moulin, 2009). We provide a mechanism that improves upon all efficient mechanisms.[8] Surprisingly, this improvement can be achieved using an extremely simple mechanism:

Corollary 1. *If the hazard rates of the type distributions are monotone, then every mechanism that allocates efficiently is dominated by a mechanism that always allocates the good to the same agent.*

More precisely, a mechanism that is better than every efficiently allocating mechanism can be found simply by comparing the agents' type distributions, giving the good to the agent with the higher expected valuation and completely ignoring any reported types.

Despite their simplicity, the optimal mechanisms obtain a surprisingly large share of first-best welfare, as the following example suggests (see Table 1 for further numerical estimates of the share of first-best welfare that the optimal mechanism obtains). Note that randomly allocating the object to one of the agents provides a worst-case welfare guarantee of $\frac{1}{2}$; in all our numerical examples the optimal posted price or option mechanism improves significantly over this lower bound.

Example 1. Suppose that $\theta_i \sim U[0, 1]$ for $i = 1, 2$. First-best welfare is given by $U_{FB} = \frac{2}{3}$, whereas the optimal mechanism M is an option mechanism with price $\frac{1}{2}$, yielding $U(M) = \frac{5}{8}$. Hence, the optimal mechanism yields a 93.8 per cent share of first-best welfare. In contrast, a random allocation yields only a 75 per cent share of first-best welfare.

The following example shows that if Condition (HR) is not satisfied the optimal mechanism need not be of the form stated in Theorem 1. The example also illustrates the role of (HR) in establishing the result.

Example 2. Let the distribution function of two symmetric agents be given as

$$f(x_i) = \begin{cases} 0.9 & \text{if } x_i \leq 0.5 \\ 0.1 & \text{otherwise.} \end{cases}$$

Due to the downwards jump at 0.5, f does not satisfy condition (HR). The optimal posted price mechanism (which is as good as the optimal option mechanism) has a strike price of $a^* = 0.275$, attaining a social welfare of 0.0718. However, the following mechanism M attains a higher social welfare of 0.0741: Set

$$d_2(x) = 1 \quad \Leftrightarrow \quad (x_2 \geq a^* \text{ and } x_1 \leq a^*) \text{ or } (x_2 \geq 0.5 \text{ and } x_1 \leq 0.5),$$

[8] Indeed, this mechanism improves upon any mechanism that treats agents symmetrically in a neighborhood of 0. This observation extends to settings with more than two agents.

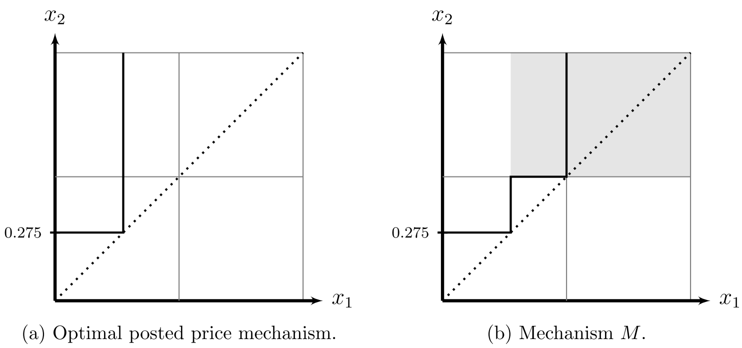

Fig. 2. Mechanisms presented in Example 2.

and set $q_2(x_1) \equiv 0$, as well as

$$q_1(x_2) = \begin{cases} 0 & \text{if } x_2 \leq a^* \\ a^* & \text{otherwise.} \end{cases}$$

This mechanism and the best option mechanism are depicted in Fig. 2. One can see that the allocation of mechanism M is more efficient. Because the induced higher payments cannot be redistributed, payments of $(0.5 - 0.275) = 0.225$ are lost for type profiles in the shaded area in Fig. 2b. But still, since type profiles x with $x_1, x_2 \geq 0.5$ appear so rarely (with density 0.01), this does not counter the positive effect due to the better allocation. In this sense, an increasing hazard rate ensures that lost payments can never be weighed out by an improved efficiency of the allocation.

5. Relaxing the constraints

In this section, we in turn relax the no-deficit, incentive compatibility and participation constraints as well as the restriction to deterministic mechanisms, and analyze how sensitive our characterization in Theorem 1 is to these relaxations.

Ex-ante budget constraints While ex-post budget constraints are imposed commonly in the literature and seem appropriate for many settings, they would effectively be turned into ex-ante constraints if insurance against budget deficits was available.[9] Relaxing the no-deficit constraint to an ex-ante constraint, thus requiring the mechanism to run at no deficit on average, simplifies the problem and allows the planner to implement the first-best (i.e., the efficient allocation and a balanced budget). This can be achieved by running the VCG mechanism. This mechanism is ex-post individually rational and creates no deficit ex-post. By redistributing the expected surplus in an arbitrary fixed way to the agents, the mechanism becomes ex-ante budget-balanced and therefore achieves the first-best.

Bayesian incentive compatible mechanisms If stronger assumptions can be made on the information structure (namely, if the agents' beliefs equal a common prior that is known to the designer), we can relax the constraints on the mechanisms to Bayesian incentive compatibility and interim individual rationality. This allows the implementation of mechanisms that achieve higher expected welfare. Notably, if the distribution of types is symmetric across agents, then the expected externality mechanism (d'Aspremont and Gerard-Varet, 1979; Arrow, 1979) achieves the first-best. To see this, observe that this mechanism allocates efficiently, has a balanced budget, and has payments given by

$$t_i^{EEM}(x) = \int_{x_i}^{\overline{x}_{-i}} x_{-i} dF_{-i}(x_{-i}) - \int_{x_{-i}}^{\overline{x}_i} x_i dF_i(x_i). \tag{4}$$

Therefore, an agent reporting a type of 0 receives a weakly positive transfer and hence a weakly positive utility. The following Proposition shows that with ex-ante symmetric agents, the expected externality mechanism is even ex-post individually rational. More generally, it shows that the first-best can be achieved whenever virtual values are increasing (in particular, under condition (HR)).

Proposition 1. *Consider the problem of finding the optimal mechanism that is Bayesian incentive compatible, interim individually rational and satisfies ex-post no-deficit.*

[9] Note also, that the exact form of the budget constraints can be irrelevant when considering Bayesian incentive compatible mechanisms (Esö and Futo, 1999).

Table 2
Simulation results comparing the welfare loss due to the restriction to deterministic mechanisms.

Distribution	Average loss	Maximum loss	Instances without loss
Random	0.018%	0.874%	92.500%
IHR	0.003%	0.420%	97.850%
Weibull	0.000%	0.000%	100.000%
All	0.007%	0.874%	96.785%

1. *If virtual valuations are increasing, (i.e., $x_i - \frac{1-F_i(x_i)}{f_i(x_i)}$ is increasing for $i = 1, 2$), then the optimal mechanism allocates efficiently and is ex-post budget-balanced.*
2. *If agents are ex-ante symmetric (i.e., $F_1 \equiv F_2$), then the expected externality mechanism is optimal. It allocates efficiently, is ex-post budget-balanced, and ex-post individually rational.*

This implies that the equivalence of Bayesian and dominant strategy incentive compatible mechanisms established by Gershkov et al. (2013) does not apply. They show that in a large class of mechanism design problems, for any Bayesian incentive compatible and interim individually rational mechanism, there exists an equivalent dominant strategy incentive compatible mechanism that is ex-post individually rational. However, this equivalence is established in the absence of budget constraints, and the above arguments imply that it cannot be extended to our setting.

Participation constraints While our general characterization of the optimal mechanism does not hold with relaxed participation constraints, these constraints are not the main driving force behind our results and the inefficiency of the optimal allocation. Indeed, our characterization can be obtained without participation constraints if one restricts attention to settings where agents are symmetric ex-ante (Shao and Zhou, 2012).

Stochastic mechanisms In the previous section we restricted attention to deterministic mechanisms in order to be able to analytically characterize the optimal mechanism. Deterministic mechanisms have additional benefits: They are simpler to implement, and more plausible in some settings (e.g., when modeling bargaining between agents).

While there are instances where the focus on deterministic mechanisms is not without loss, numerical simulations suggest that the induced loss in welfare is small. We generated $n = 2000$ random instances for three classes of distributions of types: Random distributions, random distributions with an increasing hazard rate, and distributions from the Weibull class with different shape and scale parameters such that the distribution has an increasing hazard rate. We then computed the optimal deterministic and stochastic mechanism for every instance. The results are summarized in Table 2 which shows, for each distribution class, the average and maximum welfare loss of the optimal deterministic mechanism, as a percentage of the welfare of the best stochastic mechanism. The fourth column shows the percentage of instances where there is no loss due to the restriction to deterministic mechanisms. As can be seen, instances where the deterministic constraint is binding appear only rarely. Further, even if this is the case, the percentage loss in expected welfare is very small. Note that whenever a stochastic mechanism is strictly better in our simulations, the optimal mechanism is not budget-balanced.

6. Bilateral trade

Myerson and Satterthwaite (1983) showed that one cannot implement the efficient allocation in the bilateral trade setting in an ex-post budget-balanced and interim individually rational way, and characterized the optimal mechanism satisfying these constraints. In the same environment, Hagerty and Rogerson (1987) study the set of dominant-strategy implementable mechanisms that are ex-post budget-balanced and individually rational, showing that essentially only posted price mechanisms fulfill these conditions. However, a priori it is not clear why one should restrict the search for the optimal mechanism to mechanisms with a balanced budget. After all, it is conceivable that deviating from a balanced budget could improve incentives and therefore lead to higher welfare. In fact, Schwartz and Wen (2012) show by example that relaxing budget-balancedness to a no-deficit constraint can improve upon posted price mechanisms. The result in this section shows that this holds only for stochastic mechanisms; when looking at deterministic mechanisms, the restriction to budget-balanced mechanisms does not reduce aggregate welfare.

Let the model and notation be as in Section 2, but assume now that agent 1 (called the “seller” from now on and indexed by S) is the owner of the good before participating in the mechanism (whereas agent 2 is called the “buyer” and indexed by B). By a buyer posted price mechanism we denote a posted price mechanism in which the buyer gets the object if and only if he announces a type high enough, and the seller a type that is low enough. Again, we are looking for a mechanism that maximizes the sum of the expected utilities of the agents, taking monetary transfers into account. The fact that in the bilateral trade setting the seller initially owns the good requires a stronger condition for a mechanism to be individually rational: now the outside option for a seller is to not participate in the mechanism and to keep the object. Hence, for a mechanism to be individually rational,

$$d_S(x)x_S + p_S(x) \geq x_S \qquad \text{and} \qquad d_B(x)x_B + p_B(x) \geq 0 \tag{IR'}$$

must hold for all $x \in X$.

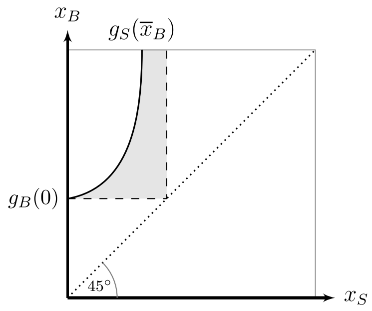

Fig. 3. Illustration of the proof of Theorem 2. The shaded area indicates the type profiles where the initial mechanism differs from the posted price mechanism with strike price $g_B(0)$ (dashed line).

Thus, a mechanism is optimal if it solves

$$\max_{M=(d,p)} \int_X \Big[d_S(x)x_S + d_B(x)x_B + p_S(x) + p_B(x) \Big] dF(x)$$
$$\text{s. t. } M \text{ satisfies IC, ND and IR'}. \tag{5}$$

Theorem 2. *There is a buyer posted price mechanism that solves problem (5).*

Proof. We first show that (IR') implies that the seller keeps the object whenever his valuation is higher. Assume to the contrary that trade takes place at $x_S > x_B$. Then (IR') for the seller implies that the seller receives at least x_S and (IR') for the buyer implies that he pays at most x_B, violating (ND).

Recall that $g_B(0)$ denotes the smallest buyer type such that trade takes place when $x_S = 0$, and $g_S(\overline{x}_B)$ denotes the highest seller type such that trade takes place when $x_B = \overline{x}_B$. We claim that $g_S(\overline{x}_B) \le g_B(0)$. Constraints (IC) and (IR') for the seller imply that he receives at least a payment of $g_S(\overline{x}_B)$ whenever the buyer reports $\overline{x}_B$ and trade takes place, in particular at $(0, \overline{x}_B)$ (if no trade takes place at $(0, \overline{x}_B)$, trade will never happen, corresponding to a posted price mechanism with a price above the highest possible valuation). Similarly, (IC) and (IR') for the buyer imply that he pays at most $g_B(0)$ whenever the seller reports 0, in particular at $(0, \overline{x}_B)$. Therefore, $g_S(\overline{x}_B) > g_B(0)$ would violate (ND) at $(0, \overline{x}_B)$.

Finally, we claim that the buyer posted price mechanism with strike price $g_B(0)$ weakly dominates the given mechanism. To see this, note that $p_S(x) + p_B(x) \le 0$ by (ND) and a posted price mechanism is budget-balanced. Hence, the posted price mechanism dominates the old mechanism with respect to payments. Since the allocation only differs for x such that $x_B \ge g_B(0) \ge x_S$ and the posted price allocation rule prescribes $d_B(x) = 1$ for such x (see also Fig. 3), the posted price mechanism also dominates the old mechanism with respect to the allocation rule. □

In contrast to Theorem 1, this result shows that a posted price mechanism is optimal for any type distribution. The difference is due to the stronger individual rationality constraint. While any allocation rule is compatible with (IR), the stronger constraint (IR') in the trade setting restricts the set of allocation rules that can be implemented without a budget deficit. Within this smaller class of feasible allocation rules, for any distribution of types a posted price mechanism is optimal.

The stronger individual rationality constraint also implies that mechanisms which do not allocate the object are infeasible. This is because if the buyer does not get the object, no money can be collected to compensate the seller for losing the object. Therefore, assuming that the object is always allocated is without loss of generality in this setting.

7. Discussion

We have studied the trade-off between efficiency and budget-balancedness in an independent private values auction model. We incorporated this into the model by letting the social welfare objective function include all payments, that is, by maximizing residual surplus.[10] We showed that, if one focuses on robust implementation in dominant strategies, an increasing hazard rate condition on agents' type distributions guarantees a resolution of the trade-off completely in favor of a balanced budget. In addition, budget-balanced mechanisms have a very simple form and can easily be implemented as posted price or option mechanisms. Further, we showed without any assumption on the prior distribution of types that a posted price mechanism is optimal in the bilateral trade setting. Our results imply that our approach of optimal mechanism design yields higher welfare than approaches concentrating on the efficient allocation.

[10] For other ways to analyze the frontier that describes possible ways to resolve the trade-off between efficiency and budget-balancedness, see, for example, Diakonikolas et al. (2012) or Tatur (2005).

In the section on robustness we have seen that the restriction to deterministic and ex-post individually rational mechanisms is not crucial for our main result. Instead, it is primarily driven by the focus on strategy-proof mechanisms that satisfy the ex-post no-deficit constraint: Without these constraints the first-best can be achieved, implying that these two restrictions are relatively costly in terms of welfare.

An interesting open question is how the result generalizes to a model including more than two agents. We strongly believe that the optimal mechanism will still be budget-balanced. An important argument for this is that, as the number of agents gets large, the efficient allocation can be approximated in a budget-balanced way: in the spirit of McAfee (1992), allocate efficiently while ignoring one agent who then receives all payments from the other agents. This can be implemented by tentatively giving the object to one of the agents and then simulating a second price auction with reserve price where this agent sells the object to the remaining agents.

Acknowledgments

We wish to thank seminar participants at the University of Bonn and Harvard University and especially Benny Moldovanu for insightful discussions. We also thank three anonymous referees and the associate editor for comments that greatly helped to improve the paper.

Appendix A. Proof of Theorem 1

The following lemma enables us to approximate any mechanism with mechanisms from the class $\mathcal{M}_0$.

Lemma 3. *For every mechanism $M = (d, p)$ and for every $\varepsilon > 0$ there exists a mechanism $\tilde{M} = (\tilde{d}, \tilde{p})$ in $\mathcal{M}_0$ such that $U(M) - U(\tilde{M}) < \varepsilon$.*

Proof. Let the mechanism $M = (d, p)$ and $\varepsilon > 0$ be given and let $g_1(x_2)$ and $g_2(x_1)$ be defined as above. Define $D_i := \{x \in X : d_i(x) = 1\}$ as the set of type profiles where agent i gets the object and define $\tilde{D}_i$ similarly. Since g_2 is a monotone function it can be approximated uniformly by a monotone and piece-wise constant function $\tilde{g}_2$. Denote the associated allocation rule by $\tilde{d}$. By choosing the step width small enough the approximation can be done such that for given $\delta > 0$,

$$\|g_1 - \tilde{g}_1\|_\infty < \delta \quad \text{and} \quad \|g_2 - \tilde{g}_2\|_\infty < \delta$$

holds. The approximation can be chosen such that $g_i(x_{-i}) = \bar{x}_i$ implies $\tilde{g}_i(x_{-i}) = \bar{x}_i$ and $\tilde{g}$ can be chosen such that $\tilde{g}_2 \leq g_2$, implying that $\tilde{D}_1 \subset D_1$.

Without loss of generality, we can assume that $q_2(x_1) \equiv 0$ (see Step 1 in the proof of Lemma 2). By construction of $\tilde{g}_2$ and since M satisfies (ND), we can define functions $\tilde{q}_i(x_{-i})$ such that $\tilde{q}_2(x_1) \equiv 0$, $0 \leq \tilde{q}_1(x_2) \leq \inf_{x_1}\{\tilde{g}_1(x_2)\tilde{d}_1(x_1, x_2) + \tilde{g}_2(x_1)\tilde{d}_2(x_1, x_2)\}$ $\forall x_2 \in X_2$ and $\|\tilde{q}_1 - q_1\|_\infty < \delta$. We then have:

$$\begin{aligned}
U(d, p) - U(\tilde{d}, \tilde{p}) &\leq \int_X q_1(x_2) - \tilde{q}_1(x_2)\, dF(x) + \int_{D_1} x_1 - g_1(x_2)\, dF(x) - \int_{\tilde{D}_1} x_1 - \tilde{g}_1(x_2)\, dF(x) \\
&\quad + \int_{D_2} x_2 - g_2(x_1)\, dF(x) - \int_{\tilde{D}_2} x_2 - \tilde{g}_2(x_1)\, dF(x) \\
&\leq \delta + \int_{D_1 \setminus \tilde{D}_1} x_1 - g_1(x_2)\, dF(x) + \int_{\tilde{D}_1} \delta\, dF(x) + \int_{\tilde{D}_2 \setminus D_2} x_2 - g_2(x_1)\, dF(x) + \int_{D_2} \delta\, dF(x) \\
&\leq 3\delta + B_1\bar{x}_1\delta + B_2\bar{x}_2\delta,
\end{aligned}$$

where B_i is an upper bound for $f_i(x_i)$. Hence, by choosing $\delta < \frac{\varepsilon}{3 + B_1\bar{x}_1 + B_2\bar{x}_2}$, it follows that $U(d, p) - U(\tilde{d}, \tilde{p}) < \varepsilon$. □

We combine the approximation lemma with Lemma 2 in order to prove the theorem.

Proof of Theorem 1. Without loss of generality, we restrict ourselves to posted price mechanisms for agent 2. We first establish that U maps the set of all posted price mechanisms to a compact subset of $\mathbb{R}$. Let $\bar{a} = \min\{\bar{x}_1, \bar{x}_2\}$ and let $a \in [0, \bar{a}]$ be some price for a posted price mechanism M_a. Then $U(M_a)$ can be written as

$$U(M_a) = \int_0^a \int_a^{\bar{x}_2} x_2 dF(x) + \int_0^{\bar{x}_1} \int_0^a x_1 dF(x) + \int_a^{\bar{x}_1} \int_a^{\bar{x}_2} x_1 dF(x).$$

Due to the continuity of F, this function is continuous with respect to a. Since $[0, \bar{a}]$ is compact, so is $\{U(M_a) \mid a \in [0, \bar{a}]\}$ and therefore there exists an a^* such that $U(M_{a^*})$ is maximal among all posted prices.

Next, assume that the theorem is false, i.e., there exists a mechanism M and $\varepsilon > 0$ such that $U(M) > U(M_{a^*}) + \varepsilon$. Then apply Lemma 3 to M and ε to get a mechanism $\tilde{M} \in \mathcal{M}_0$ with $U(\tilde{M}) > U(M_{a^*})$. This contradicts Lemma 2, establishing the theorem. $\square$

Proof of Corollary 1. The arguments in Step 1 in the proof of Lemma 2 imply that agent 1 receives no redistribution payments; symmetric arguments imply that agent 2 also gets no redistribution payments. Hence, all payments that are collected must be wasted. The result then follows from Hartline and Roughgarden (2008). $\square$

Proof of Proposition 1. (1) Let d^* denote an efficient allocation rule. Given a mechanism (d^*, t), let

$$U_i(x_i) := \int_{\underline{x}_{-i}}^{\overline{x}_{-i}} x_i \cdot d_i^*(x_i, x_{-i}) + t_i(x_i, x_{-i}) dF_{-i}(x_{-i})$$

denote the interim expected utility, and $S := -\mathbb{E}[t_1(x_1, x_2) + t_2(x_1, x_2)]$ the expected budget surplus. Observe that

$$\begin{aligned}
&\int_X x_1 \cdot d_1^*(x_1, x_2) + x_2 \cdot d_2^*(x_1, x_2) dF(x_1, x_2) - S \\
&= \int_X U_1(x_1) + U_2(x_2) dF(x_1, x_2) \\
&= U_1(0) + U_2(0) + \int_X \int_0^{x_1} d_1^*(s, x_2) ds + \int_0^{x_2} d_2^*(x_1, s) ds \, dF(x_1, x_2) \\
&= U_1(0) + U_2(0) + \int_X \frac{1 - F_1(x_1)}{f_1(x_1)} d_1^*(x_1, x_2) + \frac{1 - F_2(x_2)}{f_2(x_2)} d_2^*(x_1, x_2) \, dF(x_1, x_2).
\end{aligned}$$

Hence,

$$U_1(0) + U_2(0) + S = \int_X \left[x_1 - \frac{1 - F_1(x_1)}{f_1(x_1)}\right] d_1^*(x_1, x_2) + \left[x_2 - \frac{1 - F_2(x_2)}{f_2(x_2)}\right] d_2^*(x_1, x_2) dF(x_1, x_2).$$

We claim that

$$U_1(0) + U_2(0) + S \geq 0. \tag{6}$$

Indeed, if $U_1(0) + U_2(0) + S < 0$ were true, then

$$\int_X \left[x_i - \frac{1 - F_i(x_i)}{f_i(x_i)}\right] d_i^*(x_1, x_2) dF(x) < 0$$

would hold for some i, as $S \geq 0$ follows from no-deficit. Together with the assumption that virtual valuations are increasing, this would imply that

$$\int_X \left[x_i - \frac{1 - F_i(x_i)}{f_i(x_i)}\right] [1 - d_i^*(x_1, x_2)] dF(x) < 0.$$

Hence

$$\int_X \left[x_i - \frac{1 - F_i(x_i)}{f_i(x_i)}\right] dF(x) < 0,$$

contradicting the fact that $\int_{X_i} x_i - \frac{1 - F_i(x_i)}{f_i(x_i)} dF_i(x_i) = 0$.

Define $t_1(x) := t_1^{EEM}(x) - \int_{X_2} t_1^{EEM}(0, s) dF_2(s)$ and $t_2(x) := -t_1(x)$. The mechanism (d^*, t) is ex-post budget balanced by construction (hence, $S = 0$); it is Bayesian incentive compatible since payments differ from the payments in the expected externality mechanism only by a constant. Moreover, $U_1(0) = 0$ by construction; therefore, $U_2(0) \geq 0$ follows from (6), showing individual rationality.

(2) Optimality of the expected externality mechanism follows from the observations before Proposition 1. Ex-post individual rationality follows from the following two facts: $x_i \leq x_j$ implies $t_i^{EEM}(x) \geq 0$ by (4), and $x_i > x_j$ similarly implies $t_i^{EEM}(x) = -\int_{x_j}^{x_i} s dF_i(s) \geq -x_i$. This implies that $x_i \cdot d_i^*(x) + t_i^{EEM} \geq 0$ for all x. $\square$

References

Arrow, K., 1979. The property rights doctrine and demand revelation under incomplete information. In: Boskin, M. (Ed.), Economics and Human Welfare. Academic Press, New York.

Athanasiou, E., 2013. A solomonic solution to the problem of assigning a private indivisible good. Games Econ. Behav. 82, 369–387.

Athey, S., Miller, D.A., 2007. Efficiency in repeated trade with hidden valuations. Theoretical Econ. 2 (3), 299–354.

Cavallo, R., 2006. Optimal decision-making with minimal waste: strategyproof redistribution of vcg payments. In: Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems. AAMAS '06. ACM, New York, NY, USA, pp. 882–889.

Chakravarty, S., Kaplan, T.R., 2013. Optimal allocation without transfer payments. Games Econ. Behav. 77 (1), 1–20.

Chawla, S., Hartline, J.D., Rajan, U., Ravi, R., 2006. Bayesian optimal no-deficit mechanism design. Internet Netw. Econ. 4286, 136–148.

Condorelli, D., 2012. What money can't buy: efficient mechanism design with costly signals. Games Econ. Behav. 75 (2), 613–624.

d'Aspremont, C., Gerard-Varet, L.-A., 1979. Incentives and incomplete information. J. Public Econ. 11 (1), 25–45.

de Clippel, G., Naroditskiy, V., Polukarov, M., Greenwald, A., Jennings, N., 2014. Destroy to save. Games Econ. Behav. 86, 392–404.

Diakonikolas, I., Papadimitriou, C., Pierrakos, G., Singer, Y., 2012. Efficiency-revenue trade-offs in auctions. In: Automata, Languages, and Programming. In: Lecture Notes in Computer Science, vol. 7392, pp. 488–499.

Esö, P., Futo, G., 1999. Auction design with a risk averse seller. Econ. Letters 65 (1), 71–74.

Faltings, B., 2005. A budget-balanced, incentive-compatible scheme for social choice. In: Agent-Mediated Electronic Commerce VI. In: Lecture Notes in Artificial Intelligence, pp. 30–43.

Gershkov, A., Goeree, J.K., Kushnir, A., Moldovanu, B., Shi, X., 2013. On the equivalence of Bayesian and dominant strategy implementation. Econometrica 81 (1), 197–220.

Green, J.R., Laffont, J.-J., 1979. Incentives in Public Decision-Making. North-Holland Pub. Co., Amsterdam, NL.

Guo, M., Conitzer, V., 2009. Worst-case optimal redistribution of vcg payments in multi-unit auctions. Games Econ. Behav. 67 (1), 69–98.

Guo, M., Conitzer, V., 2010. Optimal-in-expectation redistribution mechanisms. Artif. Intell. 174 (5–6), 363–381.

Guo, M., Conitzer, V., 2014. Better redistribution with inefficient allocation in multi-unit auctions. Artif. Intell. 216, 287–308.

Hagerty, K.M., Rogerson, W.P., 1987. Robust trading mechanisms. J. Econ. Theory 42 (1), 94–107.

Hartline, J.D., Roughgarden, T., 2008. Optimal mechanism design and money burning. In: Proceedings of the 40th Annual ACM Symposium on Theory of Computing. STOC '08. ACM, pp. 75–84.

Manelli, A.M., Vincent, D.R., 2010. Bayesian and dominant-strategy implementation in the independent private-values model. Econometrica 78 (6), 1905–1938.

McAfee, R.P., 1992. A dominant strategy double auction. J. Econ. Theory 56 (2), 434–450.

McAfee, R.P., McMillan, J., 1992. Bidding rings. Amer. Econ. Rev. 82 (3), 579–599.

Miller, D., 2012. Robust collusion with private information. Rev. Econ. Stud. 79 (2), 778–811.

Moulin, H., 2009. Almost budget-balanced vcg mechanisms to assign multiple objects. J. Econ. Theory 144 (1), 96–119.

Myerson, R.B., 1981. Optimal auction design. Math. Oper. Res. 6 (1), 58–73.

Myerson, R.B., Satterthwaite, M.A., 1983. Efficient mechanisms for bilateral trading. J. Econ. Theory 29 (2), 265–281.

Schwartz, J.A., Wen, Q., 2012. Robust trading mechanisms with budget surplus and partial trade. Working paper. Kennesaw State University and Vanderbilt University.

Shao, R., Zhou, L., 2012. Optimal allocation of an indivisible good. Working paper. Yeshiva University and Shanghai Jiao Tong University.

Sprumont, Y., 2013. Constrained-optimal strategy-proof assignment: beyond the groves mechanisms. J. Econ. Theory 148, 1102–1121.

Tatur, T., 2005. On the trade off between deficit and inefficiency and the double auction with a fixed transaction fee. Econometrica 73 (2), 517–570.