

Abeliuk, Andrés; Berbeglia, Gerardo; Van Hentenryck, Pascal

Article

Bargaining mechanisms for one-way games

Games

Provided in Cooperation with:

MDPI – Multidisciplinary Digital Publishing Institute, Basel

Suggested Citation: Abeliuk, Andrés; Berbeglia, Gerardo; Van Hentenryck, Pascal (2015) : Bargaining mechanisms for one-way games, Games, ISSN 2073-4336, MDPI, Basel, Vol. 6, Iss. 3, pp. 347-367, <https://doi.org/10.3390/g6030347>

This Version is available at:

<https://hdl.handle.net/10419/167949>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<http://creativecommons.org/licenses/by/4.0/>

Article

Bargaining Mechanisms for One-Way Games *

Andrés Abeliuk ^{1,2,*}, Gerardo Berbeglia ^{1,3} and Pascal Van Hentenryck ^{1,4}

¹ National Information and Communications Technology Australia, NICTA Victoria Lab, West Melbourne, VIC 3003, Australia

² Department of Computing and Information Systems, The University of Melbourne, Parkville, VIC 3010, Australia

³ Centre for Business Analytics, Melbourne Business School, The University of Melbourne, Parkville, VIC 3010, Australia; E-Mail: g.berbeglia@mbs.edu

⁴ Department of Industrial and Operations Engineering, University of Michigan, 1205 Beal Avenue, Ann Arbor, MI 48109, USA; E-Mail: pvanhent@umich.edu

* Author to whom correspondence should be addressed; E-Mail: andres.abeliuk@nicta.com.au; Tel.: +61-403-754-676.

Academic Editor: Bahar Leventoglu

Received: 1 June 2015 / Accepted: 27 August 2015 / Published: 8 September 2015

Abstract: We introduce one-way games, a two-player framework whose distinguishable feature is that the private payoff of one (independent) player is determined only by her own strategy and does not depend on the actions taken by the other (dependent) player. We show that the equilibrium outcome in one-way games without side payments and the social cost of any *ex post* efficient mechanism can be far from the optimum. We also show that it is impossible to design a Bayes–Nash incentive-compatible mechanism for one-way games that is budget-balanced, individually rational and efficient. To address this negative result, we propose a privacy-preserving mechanism based on a single-offer bargaining made by the dependent player that leverages the intrinsic advantage of the independent player. In this setting the outside option of the dependent player is not known a priori; however, we show that the mechanism satisfies individual rationality conditions, is incentive-compatible, budget-balanced and produces an outcome that is more efficient than the equilibrium without

*An earlier, shorter version of this paper appeared in the Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI), 2015 [1].

payments. Finally, we show that a randomized multi-offer extension brings no additional benefit in terms of efficiency.

Keywords: bargaining; mechanism design; price of anarchy; distributed problem solving

1. Introduction

When modeling economic interactions between agents, it is standard to adopt a general framework where payoffs of individuals are dependent on the actions of all other decision-makers. However, some agents may have payoffs that depend only on their own actions, not on actions taken by other agents. In this paper, we explore the consequences of such asymmetries among agents. Since these features lead to a restricted version of the general model, the hope is that we can identify mechanisms that produce efficient outcomes by exploiting the properties of this specific setting.

A classic application of this setting is Coase's example of a polluter and a single victim, e.g., a steel mill that affects a laundry. The Coase theorem [2] is often interpreted as a demonstration of why private negotiations between polluters and victims can yield efficient levels of pollution without government interference. However, in an influential article, Hahnel and Sheeran [3] criticize the Coase theorem by showing that, under more realistic conditions, it is unlikely that an efficient outcome will be reached. They emphasize that the solution is a negotiation and not a market-based transaction, as described by Coase. As such, incomplete information plays an important role, and game theory and bargaining games can explain inefficient outcomes.

This paper aims at taking the first step in this direction by proposing a class of two-player one-way dependent decision settings, which formalizes many of Hahnel and Sheeran's critiques. We present a number of negative and positive results on one-way games. We first show that Nash equilibria in one-way games, under no side payments, can be arbitrarily far from the optimal social welfare. Moreover, in contrast to the Coase theorem, we show that when side payments are allowed in a Bayes–Nash incentive-compatible setting, there is no *ex post* efficient individually rational and budget-balanced mechanism for one-way games. To address this negative result, we focus on mechanisms that are budget-balanced, individually rational and incentive-compatible that are relatively efficient. Our main positive result is a single-offer bargaining mechanism, which, under reasonable assumptions for the players, increases the social welfare compared to the setting where no side payments are allowed. We also show that this single-offer mechanism cannot be improved by a (randomized) multi-offer mechanism.

This paper is an extended version of the work presented at IJCAI-15 [1], and contains additional details, examples, and several generalizations and extensions of the model. More precisely, we have generalized our framework by analyzing the case where the outside option of one of the players is not known a priori in the bargaining mechanism. We also introduce a randomized multi-offer extension of the single-offer bargaining and show that it brings no additional benefit in terms of efficiency.

The rest of this paper is organized as follows. In Section 2, we define one-way games and study the properties of Nash equilibria without payments. In Section 4.1, we prove an impossibility result.

The single-offer mechanism is presented and analyzed in Sections 5 and 5.1, respectively. Finally, in Section 6, we present a multi-offer mechanism, and we show that it does not improve the efficiency with respect to the single-offer one.

2. One-Way Games

One-way games feature two players A and B . Each player $i \in A, B$ has a public action set \mathcal{S}_i , and we write $\mathcal{S} = \mathcal{S}_A \times \mathcal{S}_B$ to denote the set of joint action profiles. As most commonly done in mechanism design, we model private information by associating each agent i with a payoff function $u_i : \mathcal{S} \times \Theta_i \rightarrow \mathbb{R}^+$, where $u_i(s, \theta_i)$ is the agent utility for strategy profile s when the agent has type θ_i . We assume that the player types are stochastically independent and drawn from a distribution f that is common knowledge. We denote by Θ_i the set of possible types of player i and write $\Theta = \Theta_A \times \Theta_B$. If $\theta \in \Theta$, we use θ_i to denote the type of player i in θ . Similar conventions are used for strategies, utilities and type distributions.

A key feature of one-way games is that the payoff $u_A((s_A, s_B), \theta_A)$ of player A is determined only by her own strategy and does not depend on B 's actions, *i.e.*,

$$\forall s_A, s_B, s'_B, \theta_A : u_A((s_A, s_B), \theta_A) = u_A((s_A, s'_B), \theta_A)$$

As a result, for the ease of notation, we use $u_A(s_A, \theta_A)$ to denote A 's payoff. Obviously, player B must act according to what player A chooses to do, and we use $s_B(s_A, \theta_B)$ to denote the best response of player B given that A plays action s_A and player B has type θ_B , *i.e.*,

$$s_B(s_A, \theta_B) = \arg\text{-max}_{s_B \in \mathcal{S}_B} u_B((s_A, s_B), \theta_B)$$

where ties are broken arbitrarily. In this paper, we always assume that ties are broken arbitrarily in arg-max expressions.

One-way games assume that players are risk-neutral agents and that after having observed the realization of their own types, players simultaneously choose their actions. As a consequence, if side payments are not allowed, player A will play an action s_A^N that yields her a maximum payoff, *i.e.*,

$$s_A^N(\theta_A) = \arg\text{-max}_{s_A \in \mathcal{S}_A} u_A(s_A, \theta_A)$$

Player B will pick $s_B^N(\theta_B)$, such that her expected payoff is maximized, *i.e.*,

$$s_B^N(\theta_B) = \arg\text{-max}_{s_B \in \mathcal{S}_B} \mathbb{E}_{\theta_A} [u_B(s_B(s_A^N(\theta_A), \theta_B), \theta_B)]$$

The set of Nash equilibria (NE) is thus characterized by $s^N(\theta) = (s_A^N(\theta_A), s_B^N(\theta_B)) \subseteq \mathcal{S}$. The best response $s_B^N(\theta_B)$ of player B may be a bad outcome for her even when B has a much greater potential payoff. Player A achieves her optimal payoff, but our motivating applications aim at optimizing a global welfare function:

$$SW((s_A, s_B), \theta) = u_A(s_A, \theta_A) + u_B((s_A, s_B), \theta_B)$$

Thus, the global welfare achieved by the Nash equilibria can be expressed as $SW(s^N(\theta), \theta)$. We quantify the quality of the Nash equilibrium outcome with the price of anarchy (PoA).

Definition 1. The price of anarchy of $s^N(\theta) \subseteq \mathcal{S}$ given type $\theta \in \Theta$ is defined as:

$$PoA(\theta) = \frac{\max_{s \in \mathcal{S}} SW(s, \theta)}{\min_{s \in s^N(\theta)} SW(s, \theta)}$$

A natural extension for the price of anarchy is to quantify the expected worst-case equilibrium.

Definition 2. The Bayes–Nash price of anarchy is defined as:

$$PoA = \mathbb{E}_\theta [PoA(\theta)]$$

Note that the price of anarchy given type $\theta \in \Theta$ can be used to obtain a lower and an upper bound on the Bayes–Nash price of anarchy in the following way,

$$\min_{\theta \in \Theta} PoA(\theta) \leq PoA \leq \max_{\theta \in \Theta} PoA(\theta)$$

Throughout this paper, we use the following two running examples to illustrate key concepts.

Example 1. Consider the instance where player A has two possible actions $s_A^1, s_A^2 \in \mathcal{S}_A$. Player’s A type $\theta_A \in [0, 100]$ is drawn from a uniform distribution between zero and 100. Action s_A^1 has a payoff $u_A(s_A^1, \theta_A) = \theta_A$, while action s_A^2 has a constant payoff $u_A(s_A^2, \theta_A) = 100$. Let player B have only one possible type θ_B . The set of dominant actions for player B corresponds to the set of best responses $s_B(s_A^1, \theta_B)$ and $s_B(s_A^2, \theta_B)$, and we set payoffs to be $u_B(s_B(s_A^1, \theta_B)) = x$ and $u_B(s_B(s_A^2, \theta_B)) = 0$, where x is a positive constant. When no transfers are allowed, player A will always play action s_A^2 , yielding a social welfare of $100 + 0$. If player A plays s_A^1 , her expected payoff is 50, and the expected social welfare is thus $50 + x$. The price of anarchy is $\frac{50+x}{100}$ if $x \geq 50$ and one otherwise. Notice that the PoA is an increasing function of x .

Example 2. Consider the instance where player A has n possible actions $s_A^1, s_A^2, \dots, s_A^n \in \mathcal{S}_A$. Player’s A type $\theta_A \in [0, 1]^n$ is drawn from n independent and identically distributed uniform distributions between zero and one. Each action has a payoff $u_A(s_A^i, \theta_A) = \theta_{A,i}$, where $\theta_{A,i}$ denotes the i -th element of vector θ_A . For player B , consider the set of best responses $s_B(s_A^1, \theta_B), s_B(s_A^2, \theta_B), \dots, s_B(s_A^n, \theta_B)$. Let the expected payoffs be $\mathbb{E}_{\theta_B} [u_B(s_B(s_A^i, \theta_B), \theta_B)] = \mu_i$ for $i = 1, 2, \dots, n$, where $\mu_1 \geq \mu_2 \geq \dots \geq \mu_n$. All other payoffs are set to zero, i.e., for any $\theta_B \in \Theta_B$, $s_A \in \mathcal{S}_A$ and $s_B \neq s_B(s_A, \theta_B)$: $u_B((s_B, s_A), \theta_B) = 0$.

When no transfers are allowed, player A chooses her (realized) maximizing payoff action. Such an action is distributed according to the largest order statistic, i.e., the maximum between all payoffs. The largest order statistic between n standard uniforms follows a $Beta(n, 1)$ distribution [4], with mean $\frac{n}{n+1}$. Hence, the expected payoff of player A is $\frac{n}{n+1}$.

By symmetry of player A ’s payoff functions, all of her actions will be played with probability $\frac{1}{n}$. Thus, player B ’s expected payoff is maximized when choosing $s_B(s_A^1, \theta_B)$, yielding her an expected payoff of $\frac{\mu_1}{n}$. As a result, the expected social welfare in equilibrium is $\frac{n}{n+1} + \frac{\mu_1}{n}$.

To compute the optimal social welfare, note that if player A select s_A^1 with probability one, her expected payoff is $\frac{1}{2}$, and the expected payoff of player B is μ_1 ; and so, the social welfare is $\frac{1}{2} + \mu_1$, which is greater than $\frac{n}{n+1} + \frac{\mu_1}{n}$ if $\frac{1}{2} + \mu_1 \geq 1 + \frac{\mu_1}{n}$ or, equivalently, $\mu_1 \geq 1$ for every value of $n \geq 2$. The

price of anarchy is $(\frac{1}{2} + \mu_1) / (\frac{n}{n+1} + \frac{\mu_1}{n})$ when $\mu_1 \geq 1$, and in the limit as $n \rightarrow \infty$, the PoA becomes $\frac{1}{2} + \mu_1$. Notice that the PoA is an increasing function of μ_1 .

Examples 1 and 2 illustrate that, when no transfers are allowed, the price of anarchy can be arbitrarily large as player B’s payoff is large compared to the payoff of player A. We now generalize this idea to quantify the price of anarchy in one-way games. Notice that, for this argument, private information plays no role. For instance, in Example 1, even if player A’s revealed type is set to be public knowledge, this would not change her strategy. However, private information will have an important role when we explore side payments in Section 5.

Proposition 1. *In one-way games, the price of anarchy when no payments are allowed satisfies, for any type θ ,*

$$\frac{\max_{s \in \mathcal{S}} u_B(s, \theta_B)}{\max_{s \in \mathcal{S}} u_A(s, \theta_A) + u_B(s^N(\theta), \theta_B)} \leq PoA(\theta) \leq 1 + \frac{\max_{s \in \mathcal{S}} u_B(s, \theta_B)}{\max_{s \in \mathcal{S}} u_A(s, \theta_A)}$$

Proof. Let $\bar{u}_i(\theta_i) = \max_{s \in \mathcal{S}} u_i(s, \theta_i)$, $i \in \{A, B\}$. The independence of player A implies that, for all $\theta \in \Theta$, her payoff is $u_A(s^N(\theta), \theta_A) = \bar{u}_A(\theta_A)$. It follows that:

$$\begin{aligned} \frac{\max\{\bar{u}_A(\theta_A), \bar{u}_B(\theta_B)\}}{\bar{u}_A(\theta_A) + u_B(s^N(\theta), \theta_B)} &\leq PoA(\theta) \\ &\leq \frac{\bar{u}_A(\theta_A) + \bar{u}_B(\theta_B)}{\bar{u}_A(\theta_A) + u_B(s^N(\theta), \theta_B)} \leq \frac{\bar{u}_A(\theta_A) + \bar{u}_B(\theta_B)}{\bar{u}_A(\theta_A)} = 1 + \frac{\bar{u}_B(\theta_B)}{\bar{u}_A(\theta_A)} \end{aligned}$$

□

The price of anarchy can thus be arbitrarily large. When it is large enough, Proposition 1 indicates that $\max_{s \in \mathcal{S}} u_B(s, \theta_B) \geq \max_{s \in \mathcal{S}} u_A(s, \theta_A) \geq u_B(s^N(\theta), \theta_B)$. In this case, player B has bargaining power to incentivize player A monetarily, so that she moves from her equilibrium and cooperates to overcome bad social welfare. This paper explores this possibility by analyzing the social welfare when side payments are allowed.

3. Related Work

Before moving to the main results, it is useful to discuss related games. One-way games may seem to resemble Stackelberg games with their notions of leader and follower. The key difference, however, is that, in one-way games, the leader does not depend on the action taken by the follower. In addition, in one-way games, players do not have complete information, and moves are simultaneous. Jackson and Wilkie [5] studied one-way instances derived from their more general framework of endogenous games. However, they tackled the problem from a different perspective and assumed complete information (*i.e.*, the player utilities are not private). Jackson and Wilkie gave a characterization of the outcome when players make binding offers of side payments, deriving the conditions under which a new outcome becomes a Nash equilibrium or remains one. They analyzed a subclass, called ‘one-sided externality’, which is essentially a one-way game, but with complete information. They showed that the efficient outcome is an equilibrium in this setting, supporting Coase’s claim that a polluter and his victim can reach an efficient outcome. Under perfect information, the victim can determine the minimal transfer necessary to support the efficient play. Naturally, this result does not hold under incomplete information [6]. In

what follows, we design a bargaining mechanism that is able to cope with the incomplete information setting.

4. Bayesian–Nash Mechanisms

In this section, we consider a Bayesian Nash setting with quasi-linear preferences. Both players A and B have private utilities and beliefs about the utilities of the other players. By the revelation principle, we can restrict our attention to direct mechanisms that implement a social choice function. A social choice function in quasi-linear environments takes the form of $f(\theta) = (k(\theta), t(\theta))$, where, for every $\theta \in \Theta$, $k(\theta) \in \mathcal{S}$ is the allocation function and $t_i(\theta) \in \mathbb{R}$ represents a monetary transfer to agent i . The main objective of mechanism design is to implement a social choice function that achieves near efficient allocations, while respecting some desirable properties. For completeness, we specify these key properties.

Definition 3. A social choice function is *ex post efficient* if, for all $\theta \in \Theta$, we have: $k(\theta) \in \arg\text{-max}_{s \in \mathcal{S}} \sum_i u_i(s, \theta)$.

Definition 4. A social choice function is *budget-balanced (BB)* if, for all $\theta \in \Theta$, we have $\sum_i t_i(\theta) = 0$.

In other words, there are no net transfers out of the system or into the system. Taken together, *ex post* efficiency and budget-balance imply Pareto optimality. An essential condition of any mechanism is to guarantee that agents report their true types. The following property captures this notion when agents have prior beliefs on the types of other agents.

Definition 5. A social choice function is *Bayes–Nash incentive compatible (IC)* if for every player i :

$$\mathbb{E}_{\theta_{-i}|\theta_i} [u_i(k(\theta_i, \theta_{-i}), \theta_i) + t_i(\theta_i, \theta_{-i})] \geq \mathbb{E}_{\theta_{-i}|\theta_i} [u_i(k(\hat{\theta}_i, \theta_i), \theta_i) + t_i(\hat{\theta}_i, \theta_{-i})]$$

where $\theta_i \in \Theta_i$ is the type of player i , $\hat{\theta}_i$ is the type player i reports and $\mathbb{E}_{\theta_{-i}|\theta_i}$ denotes player i 's expectation over prior beliefs θ_{-i} of the types of other agents given her own type θ_i .

The most natural definition of individual rationality (IR) is interim IR, which states that every agent type has non-negative expected gains from participation.

Definition 6. A social choice function is *interim individual rational* if, for all types $\theta \in \Theta$, it satisfies:

$$\mathbb{E}_{\theta_{-i}|\theta_i} [u_i(k(\theta), \theta_i) + t_i(\theta)] \geq \bar{u}_i(\theta_i)$$

where $\bar{u}_i(\theta_i)$ is the expected utility for non-participation.

In the context of one-way games, both players have positive outside options that depend only on their types. In particular, the outside options are given by the Nash equilibrium outcome under no side payments. For players A and B , the expected utilities for non-participation are $\bar{u}_A(\theta_A) = u_A(s_A^N(\theta_A), \theta_A)$ and $\bar{u}_B(\theta_B) = u_B(s^N(\theta), \theta_B)$, respectively.

4.1. Impossibility Result

This section shows that there exists no mechanism for one-way games that is efficient and satisfies the traditional desirable properties. The result is derived from the Myerson–Satterthwaite [6] theorem, a seminal impossibility result in mechanism design. The Myerson–Satterthwaite theorem considers a bargaining game with two-sided private information, and it states that, for a bilateral trade setting, there exists no Bayes–Nash incentive-compatible mechanism that is budget balanced, *ex post* efficient and that gives every agent type non-negative expected gains from participation (*i.e.*, *ex interim* individual rationality).

Our contribution is two-fold: we present an impossibility result for one-way games, and we relate them with bargaining games, an idea that we will further explore in the following sections. We now formalize the impossibility result for one-way games.

Consider the Myerson–Satterthwaite bilateral bargaining setting.

Definition 7. *Myerson–Satterthwaite bargaining game:*

1. A seller (Player 1) owns an object for which her valuation is $v_1 \in V_1$, and a buyer (Player 2) wants to buy the object at a valuation $v_2 \in V_2$.
2. Each player i knows her valuation v_i at the time of the bargaining, and Player 1 (respectively, 2) has a probability density distribution $f_2(v_2)$ (respectively, $f_1(v_1)$) for the other player's valuation.
3. Both distributions are assumed to be continuous and positive on their domain, and the intersection of the domains is not empty.

By the revelation principle, we can restrict our attention to incentive-compatible direct mechanisms. A direct mechanism for bargaining games is characterized by two functions: (1) a probability distribution $\sigma : V_1 \times V_2 \rightarrow [0, 1]$ that specifies the probability that the object is transferred from the seller to the buyer; and (2) a monetary transfer scheme $p : V_1 \times V_2 \rightarrow \mathbb{R}^2$. In this setting, *ex post* efficiency is achieved if $\sigma(v_1, v_2) = 1$ when $v_1 < v_2$, and zero otherwise.

Our result consists of showing that a mechanism \mathcal{M}' for the Myerson–Satterthwaite setting can be constructed using a mechanism \mathcal{M} for a one-way game in such a way that, if \mathcal{M} is efficient, individual rational (IR), incentive compatible (IC) and budget-balanced (BB), then \mathcal{M}' is efficient, IR, IC and BB. The Myerson–Satterthwaite impossibility theorem states that such a mechanism \mathcal{M}' cannot exist, which implies the following impossibility result for one-way games.

Theorem 1. *There is no ex post efficient, individually rational, incentive-compatible and budget-balanced mechanism for one-way games.*

Proof. For any bargaining setting, consider the following transformation into a one-way game instance:

$$\mathcal{S}_A = \{s_A^1, s_A^2\}, \mathcal{S}_B = \{s_B\}$$

$$\forall v_1 \in V_1 : u_A(s_A^1, v_1) = v_1, u_A(s_A^2, v_1) = 0$$

$$\forall v_2 \in V_2 : u_B((s_A^1, s_B), v_2) = 0, u_B((s_A^2, s_B), v_2) = v_2$$

where player types $(v_1, v_2) \in V_1 \times V_2$ are drawn from distribution $f_1 \times f_2$. Two possible outcomes may occur, (s_A^1, s_B) or (s_A^2, s_B) , with social welfare v_1 and v_2 , respectively.

Let us assume $\mathcal{M} = (k, t)$ is a direct mechanism for one-way games and that \mathcal{M} is *ex post* efficient, IR, IC and BB. We now construct a mechanism $\mathcal{M}' = (\sigma, p)$, where $\sigma(v_1, v_2)$ is the probability that the object is transferred from the seller to the buyer and $p(v_1, v_2)$ is the payment of each player. We define \mathcal{M}' such that:

$$\sigma(v_1, v_2) = \begin{cases} 0 & \text{if } k(v_1, v_2) = (s_A^1, s_B) \\ 1 & \text{if } k(v_1, v_2) = (s_A^2, s_B) \end{cases}$$

and:

$$p(v_1, v_2) = t(v_1, v_2)$$

It remains to show that \mathcal{M}' satisfies all of the desired properties. An *ex post* efficient mechanism \mathcal{M} in the one-way instance satisfies:

$$k(v_1, v_2) = \begin{cases} (s_A^1, s_B) & \text{if } v_1 \geq v_2 \\ (s_A^2, s_B) & \text{if } v_1 < v_2 \end{cases}$$

Therefore, $\sigma(v_1, v_2)$ will assign the object to the buyer iff $v_1 < v_2$. That is, the player with the highest valuation will always get the object, meeting the restriction of *ex post* efficiency. The budget-balanced constraint in \mathcal{M} implies that $p_1(v_1, v_2) + p_2(v_1, v_2) = 0$ for all possible valuations, so \mathcal{M}' is budget-balanced.

The individual rationality property for \mathcal{M}' comes from noticing that the default strategy of player A when no payments are allowed is s_A^1 , and the corresponding payoff is v_1 . Therefore, the seller utility is guaranteed to be at least her valuation v_1 . Analogously, the buyer will not have a negative utility given that $u_B((s_A^1, s_B), v_2) = 0$.

Incentive compatibility is straightforward from the definition. Assume that \mathcal{M}' is not incentive compatible, then in mechanism \mathcal{M} , at least one player could benefit from reporting a false type.

Such a mechanism \mathcal{M}' cannot exist, since it contradicts the Myerson–Satterthwaite impossibility result, which concludes our proof. \square

An immediate consequence of this result is that Bayesian Nash mechanisms can only achieve at most two of the three properties: *ex post* efficiency, individual rationality and budget balance. For instance, the well-known Vickrey–Clarke–Groves (VCG) mechanism and the d’Aspremont, Gérard-Varet [7] and Arrow [8] (dAGVA) mechanism are part of the Groves family of mechanisms that truthfully implement social choice functions that are *ex post* efficient. VCG has no guarantee of budget balance, while dAGVA is not guaranteed to meet the individual rationality constraints. We refer the reader to Williams [9] and Krishna and Perry [10] for alternative derivations of the impossibility result for bilateral trading under the Groves family of mechanisms.

5. Single-Offer Mechanism

In this section, we propose a simple bargaining mechanism for player B to increase her payoff. The literature about bargaining games is extensive, and we refer readers to a broad review by Kennan and Wilson [11].

Given the nature of our applications, individual rationality imposes a necessary constraint. Otherwise, player A can always defect from participating in the mechanism and achieve her maximal payoff independently of the type of player B . Additionally, we search for Bayesian Nash mechanisms without subsidies, *i.e.*, budget-balanced mechanisms. The lack of a subsidiary in this case gives rise to a decentralized mechanism that does not require a third agent to perform the computations needed by the mechanism. However, a third party is needed to ensure compliance with the agreement reached by both players.

An interesting starting point for one-way games is the recognition that, whenever player B has a better payoff than A , player A may let player B play her optimal strategy in exchange for money. The resulting outcome can be viewed as swapping the roles of both players, *i.e.*, player B chooses her optimal strategy and A plays her best response to B 's strategy. In this case, as in Proposition 1, the worst outcome would be:

$$1 + \frac{\max_{s \in \mathcal{S}} u_A(s, \theta_A)}{\max_{s \in \mathcal{S}} u_B(s, \theta_B)}$$

This observation together with Proposition 1 leads to the following lemma.

Lemma 1. *Consider the social choice function that selects the best action that maximizes the payoff of either player A or player B , *i.e.*, the strategy:*

$$s'(\theta) = \arg\text{-max}_{s \in \mathcal{S}} (\max(u_A(s, \theta_A), u_B(s, \theta_B)))$$

*In the one-way game, the price of anarchy of implementing strategy $s'(\theta)$ is two (*i.e.*, $\forall \theta \text{ } P_oA(\theta) = 2$).*

The intuition behind this result is that the optimal social welfare is upper bounded by twice the utility of the action that maximizes the payoff of either player, whereas in the worst case scenario, the payoff of one of the players may be zero.

Unfortunately, this social choice function cannot be implemented in dominant strategies without violating individual rationality. Player A may have a smaller payoff by following strategy s' instead of the Nash equilibrium strategy s^N . Indeed, when $SW(s', \theta) < SW(s^N, \theta)$, for any budget-balanced transfer function, it must be that at least one of the players will be worse than playing the Nash equilibrium strategy s^N . Lemma 1, however, gives us hope for designing a budget-balanced mechanism that has a constant price of anarchy. Indeed, a simple and distributed implementation would ask player B to propose an action to be implemented, and player A would receive a monetary compensation for deviating from her maximal strategy.

We now present such a distributed implementation based on a bargaining mechanism. The mechanism is inspired by the model of two-person bargaining under incomplete information presented by Chatterjee and Samuelson [12]. In their model, both the seller and the buyer submit sealed offers, and a trade occurs if there is a gap in the bids. The price is then set to be a convex combination of the bids. Our single-offer mechanism adapts this idea to one-way games. In particular, to counteract player A 's advantage, player B makes the first and final offer. Moreover, the structure of our mechanism makes it possible to quantify the price of anarchy and to provide a quality guarantee on the mechanism outcome. Our single-offer mechanism is defined as follows:

1. Player B selects an action $s_A \in \mathcal{S}_A$ to propose to player A .
2. Player B also computes her outside option $s_B^O(s_A, \theta_B)$ in case player A rejects action s_A , and we denote by $u_B^O(s_A, \theta_B)$ the expected payoff from her outside option.
3. Player B proposes a monetary value of $\gamma \cdot \delta_B(s_A, \theta_B)$ with $\delta_B(s_A, \theta_B) = u_B(s_B(s_A, \theta_B), \theta_B) - u_B^O(s_A, \theta_B)$ and $\gamma \in \mathbb{R}_{[0,1]}$ to player A in the hope that she accepts to play strategy s_A instead of strategy s_A^N .
4. Player A decides whether to accept the offer.
5. If player A accepts the offer, the outcome of the game is $(s_A, s_B(s_A, \theta_B))$. Otherwise, the outcome of the game is the outside option $(s_A^N(\theta_A), s_B^O(s_A, \theta_B))$.

It is worth observing that a broker is required in this mechanism to ensure that the outcome $(s_A^N(\theta_A), s_B^O(s_A, \theta_B))$ is implemented if player A rejects the unique offer, and no counteroffers are made. A key feature of the single-offer mechanism is that it requires a minimum amount of information from player A (i.e., whether she accepts or rejects the offer).

To derive the equilibrium strategy for the single-offer mechanism, we assume that players are expected utility maximizers. The parameter $\gamma \in \mathbb{R}_{[0,1]}$ has been chosen so that player B , satisfying individual rationality, never offers more than $\delta_B(s_A, \theta_B)$, and her payoff is never worse than her expected outside option $u_B^O(s_A, \theta_B)$. Whereas the mechanism can only guarantee interim individual rationality for player B , it provides *ex post* individual rationality for player A , as shown in the following proposition.

Proposition 2. *If players A and B play the single-offer mechanism, for any $(\theta_A, \theta_B) \in \Theta$, player A accepts the offer (γ, s_A) whenever:*

$$u_A(s_A, \theta_A) + \gamma \cdot \delta_B(s_A, \theta_B) \geq u_A(s_A^N(\theta_A), \theta_A)$$

In case player A rejects the offer (γ, s_A) , she will choose her utility maximizing action $s_A^N(\theta_A)$ as her outside option. Note that by Proposition 2, if $s_A = s_A^N(\theta_A)$, player A would never reject the proposed action s_A . Accordingly, if proposed action s_A is rejected, then $s_A \neq s_A^N(\theta_A)$. This observation leads to the following proposition.

Proposition 3. *For every task $s \in \mathcal{S}_A$, let $\Theta_A^s = \Theta_A \setminus \{\theta_A \in \Theta_A : s = \arg\text{-max}_{x \in \mathcal{S}_A} u_A(x, \theta_A)\}$. In the single-offer mechanism, player B will pick outside option $s_B^O(s_A, \theta_B)$ such that her expected payoff is maximized, i.e.,*

$$s_B^O(s_A, \theta_B) = \arg\text{-max}_{s_B \in \mathcal{S}_B} \mathbb{E}_{\theta_A^s} [u_B((s_A^N(\theta_A), s_B), \theta_B)]$$

Example 1. (continued) The payoff of player B is higher if action s_A^1 is played by player A . Hence, player B has incentives to submit an offer c that triggers action s_A^1 . Player A of type θ_A accepts the offer if $c + u_A(s_A^1, \theta_A) \geq u_A(s_A^2, \theta_A) = 100$. Given that $u_A(s_A^1, \theta_A)$ follows a uniform distribution, the probability that player A accepts the offer is $\frac{c}{100}$ if $c \leq 100$ and one otherwise. For player B , this offer has an expected payoff of $\frac{c}{100} \cdot (x - c)$ if $c \leq 100$ and $x - c$ otherwise. The optimal value for the offer

is given by $c^* = \frac{x}{2}$ if $x \leq 200$ and $c^* = 100$ if $x > 200$. This leads to an expected social welfare for the single-offer mechanism of:

$$SW = \begin{cases} 100 + x \left(\frac{x}{200} - \frac{1}{4} \right) & \text{if } x \leq 200 \\ 50 + x & \text{if } x > 200 \end{cases}$$

Recall that the optimal social welfare is $50+x$ if $x \geq 50$ and 100 otherwise. Therefore, the mechanism has a price of anarchy,

$$PoA = \begin{cases} \frac{100}{100+x \left(\frac{x}{200} - \frac{1}{4} \right)} & \text{if } x \leq 50 \\ \frac{50+x}{100+x \left(\frac{x}{200} - \frac{1}{4} \right)} & \text{if } 50 \leq x \leq 200 \\ \frac{50+x}{50+x} & \text{if } 200 \leq x \end{cases}$$

The PoA is bounded by a constant, and in fact, $PoA \leq 1.21$ for any x . This contrasts with the unbounded PoA obtained when no side payments are allowed.

Example 2. (continued) The payoff of player B is higher if action s_A^1 is played by player A . Hence, player B has an incentive to submit a monetary offer $c \leq 1$ that triggers action s_A^1 . Player A of type θ_A accepts the offer if $c + u_A(s_A^1, \theta_A) \geq \max_{s_A \in S_A} u_A(s_A, \theta_A)$. It can be shown that the probability that player A accepts the offer is $\frac{cn - c^n}{n - 1}$. In case of acceptance, the expected payoff is $\mu_1 - c$ for player B and $\frac{1}{2} + c$ for player A . In case of rejection, it is guaranteed that player A will not play s_A^1 , and hence, player B 's outside option is action $s_B(s_A^2, \theta_B)$ with an expected payoff of $\frac{\mu_2}{n-1}$. Player A 's expected outside option is $\frac{n}{n+1}$, corresponding to her expected maximum payoff derived in Example 2. As a result, player B by offering c , has an expected payoff of:

$$\frac{cn - c^n}{n - 1} (\mu_1 - c) + \left(1 - \frac{cn - c^n}{n - 1} \right) \frac{\mu_2}{n + 1}$$

When n is large, the probability of acceptance is approximately,

$$\lim_{n \rightarrow \infty} \frac{cn - c^n}{n - 1} = c$$

Accordingly, in case of rejection, the expected payoffs become $\lim_{n \rightarrow \infty} \frac{\mu_2}{n-1} = 0$ for player B and $\lim_{n \rightarrow \infty} \frac{n}{n+1} = 1$ for player A . Player B 's expected payoff is thus $c(\mu_1 - c)$ and is maximized when she offers $c^* = \frac{\mu_1}{2}$ if $\mu_1 \leq 2$ and $c^* = 1$ otherwise. This leads to an expected social welfare for the single-offer mechanism of $SW = c^*(\mu_1 + \frac{1}{2}) + (1 - c^*)(0 + 1)$. Recall that the optimal social welfare is $\frac{1}{2} + \mu_1$ if $\mu_1 \geq \frac{1}{2}$ and one otherwise. Therefore, the mechanism has the following price of anarchy,

$$PoA = \begin{cases} \frac{1}{\frac{\mu_1}{2}(\mu_1 + \frac{1}{2}) + (1 - \frac{\mu_1}{2})} & \text{if } \mu_1 \leq \frac{1}{2} \\ \frac{\frac{1}{2} + \mu_1}{\frac{\mu_1}{2}(\mu_1 + \frac{1}{2}) + (1 - \frac{\mu_1}{2})} & \text{if } \frac{1}{2} \leq \mu_1 \leq 2 \\ \frac{\frac{1}{2} + \mu_1}{\frac{1}{2} + \mu_1} = 1 & \text{if } 2 \leq \mu_1 \end{cases}$$

and the PoA has a maximum value of $\frac{4}{31} (3 + 2\sqrt{10}) \approx 1.203$. This contrasts with the unbounded PoA obtained by the Nash equilibrium when no side payments are allowed.

We now generalize the analysis done in Examples 1 and 2. We proceed by studying the utility-maximizing strategy (s_A, γ) for player B and then derive the expected social welfare of the outcome for the single-offer mechanism. Note that, in case of agreement, the action of player B of type θ_B is solely defined by s_A , as she has no incentives to defect from its best response $s_B(s_A, \theta_B)$. By Proposition 2, player A accepts an offer whenever $\delta_A(s_A, \theta_A) \leq \gamma \cdot \delta_B(s_A, \theta_B)$, where $\delta_A(s_A, \theta_A) = u_A(s_A^N(\theta_A), \theta_A) - u_A(s_A, \theta_A)$. Player B obviously aims at choosing γ and s_A to maximize her payoff, and we now study this optimization problem. In the case of an agreement, player B is left with a profit of:

$$u_B(s_B(s_A, \theta_B), \theta_B) - \gamma \cdot \delta_B(s_A, \theta_B)$$

Otherwise, player B gets an expected payoff of $u_B^O(s_A, \theta_B)$.

Definition 8. The probability that player A accepts the offer (s_A, γ) , given that player B has type $\theta_B \in \Theta_B$, is:

$$P(s_A, \gamma, \theta_B) = \Pr[\gamma \cdot \delta_B(s_A, \theta_B) \geq \delta_A(s_A, \theta_A)] = \int_{\theta_A \in \Theta_A} f_A(\theta_A) \cdot \mathbb{1}(s_A, \gamma \cdot \delta_B(s_A, \theta_B), \theta_A) d\theta_A$$

with:

$$\mathbb{1}(s_A, x, \theta_A) = \begin{cases} 1 & \text{if } x \geq \delta_A(s_A, \theta_A) \\ 0 & \text{otherwise} \end{cases}$$

The expected profit of players A and B for proposed action $s = (s_A, s_B)$ and γ when player B has type θ_B is given by:

$$\begin{aligned} \mathbb{E}_{\theta_A}[U_B(s_A, \gamma, \theta_B)] &= u_B^O(s_A, \theta_B) + P(s_A, \gamma, \theta_B) \cdot ((1 - \gamma) \cdot \delta_B(s_A, \theta_B)) \\ \mathbb{E}_{\theta_A}[U_A(s_A, \gamma, \theta_B)] &= \mathbb{E}_{\theta_A}[u_A^N(\theta_A)] + P(s_A, \gamma, \theta_B) \cdot (\gamma \cdot \delta_B(s_A, \theta_B) - \mathbb{E}_{\theta_A}[\delta_A(s_A, \theta_A)]) \end{aligned}$$

The optimal strategy of player B is specified in the following lemma.

Lemma 2. On the single-offer mechanism, player B chooses $s_A^*(\theta_B)$ and $\gamma^*(s_A^*, \theta_B)$ such that:

$$s_A^*(\theta_B) = \arg\text{-max}_{s_A \in \mathcal{S}_A} \mathbb{E}_{\theta_A}[U_B(s_A, \gamma^*, \theta_B)]$$

where:

$$\gamma^*(s_A, \theta_B) = \arg\text{-max}_{\gamma \in \mathbb{R}_{[0,1]}} P(s_A, \gamma, \theta_B) \cdot (1 - \gamma)$$

5.1. Price of Anarchy

We now analyze the quality of the outcomes in the single-offer mechanism. The first step is the derivation of a lower bound for the expected social welfare of the single-offer mechanism. Inspired by Lemma 1, instead of considering all pairs $\langle s_A, \gamma \rangle$, the analysis restricts attention to a single action $s'_A = \arg\text{-max}_{s_A \in \mathcal{S}_A} u_B(s_B(s_A, \theta_B), \theta_B)$. We prove that, when offering to player A action s'_A and its associated optimal value for γ , the expected social welfare is lower than the optimal pair $\langle s_A^*, \gamma^* \rangle$. As a result, we obtain an upper bound to the price of anarchy of the single-offer mechanism.

To make the discussion precise, consider the strategy where player B offers $\langle s'_A, \gamma^*(s'_A, \theta_B) \rangle$, with $\gamma^*(s'_A, \theta_B)$ being the optimal choice of γ given s'_A , following the notation used in Lemma 2.

Lemma 3. For any type $\theta_B \in \Theta_B$ of player B , the expected social welfare achieved by the single-offer mechanism is at least the expected social welfare achieved by the strategy $\langle s'_A, \gamma^*(s'_A, \theta_B) \rangle$.

Proof. Let $\gamma^* = \gamma^*(s_A^*, \theta_B)$ and $\gamma' = \gamma^*(s'_A, \theta_B)$. The optimality condition of s^* implies that:

$$\mathbb{E}_{\theta_A} [U_B(s', \gamma', \theta_B)] \leq \mathbb{E}_{\theta_A} [U_B(s^*, \gamma^*, \theta_B)] \tag{1}$$

Two cases can occur. The first case is:

$$P(s'_A, \gamma', \theta_B) \leq P(s_A^*, \gamma^*, \theta_B)$$

i.e., the probability of player A accepting offer (s_A^*, γ^*) is greater than if offered (s'_A, γ') . Then, it must be that the expected payoff of player A is greater when offered (s_A^*, γ^*) , *i.e.*,

$$\mathbb{E}_{\theta_A} [U_A(s'_A, \gamma', \theta_B)] \leq \mathbb{E}_{\theta_A} [U_A(s_A^*, \gamma^*, \theta_B)]$$

This, together with Inequality (1), results in the single-offer mechanism having a greater expected social welfare.

The second case is:

$$P(s'_A, \gamma', \theta_B) > P(s_A^*, \gamma^*, \theta_B)$$

Consider γ'' such that $P(s'_A, \gamma'', \theta_B) = P(s_A^*, \gamma^*, \theta_B)$. The fact that the probabilities of acceptance are the same implies that the expected payoff of player A is the same in both cases, *i.e.*, $\mathbb{E}_{\theta_A} [U_A(s'_A, \gamma'', \theta_B)] = \mathbb{E}_{\theta_A} [U_A(s_A^*, \gamma^*, \theta_B)]$. This, together with Equation (1), yields:

$$\mathbb{E}_{\theta_A} [SW(s^*, \gamma^*, \theta_B)] \geq \mathbb{E}_{\theta_A} [SW(s', \gamma'', \theta_B)]$$

This is equivalent to:

$$u_B(s^*, \theta_B) + \mathbb{E}_{\theta_A} [u_A(s_A^*, \theta_A)] \geq u_B(s', \theta_B) + \mathbb{E}_{\theta_A} [u_A(s'_A, \theta_A)] \tag{2}$$

Similarly, consider γ^{**} such that:

$$P(s'_A, \gamma', \theta_B) = P(s_A^*, \gamma^{**}, \theta_B)$$

which implies:

$$\mathbb{E}_{\theta_A} [U_A(s'_A, \gamma', \theta_B)] = \mathbb{E}_{\theta_A} [U_A(s_A^*, \gamma^{**}, \theta_B)]$$

The existence of γ^{**} is guaranteed by Inequality (2), which states that there is more money in expectation to transfer to player A when choosing s^* over s' . The fact that the acceptance probabilities are the same, together with Inequality (2), implies that:

$$\mathbb{E}_{\theta_A} [SW(s^*, \gamma^{**}, \theta_B)] \geq \mathbb{E}_{\theta_A} [SW(s', \gamma', \theta_B)]$$

Given that the expected payoff of player A is the same in both cases, it must be the case that the expected payoff of player B is higher when using (s_A^*, γ^{**}) .

Therefore, we have found an offer for the single-offer mechanism with greater expected social welfare and a greater payoff for player B compared to strategy $\langle s'_A, \gamma' \rangle$. \square

We are ready to derive an upper bound for the induced price of anarchy for the single-offer mechanism. We first derive the price of anarchy of strategy $\langle s'_A, \gamma' \rangle$ in the case of agreement and disagreement of player A .

Lemma 4. Consider action $s' = \arg \max_{s \in S} u_B(s, \theta_B)$, and let $PoA^A(\gamma)$ and $PoA^R(\gamma)$ denote the induced price of anarchy if player A accepts and rejects the offer given a proposed γ . Then,

$$PoA^A(\gamma) = 1 + \gamma \quad \text{and} \quad PoA^R(\gamma) = 1 + \frac{1}{\gamma}$$

Proof. Let $s_A^N = s_A^N(\theta_A)$, $u_A^N = u_A(s_A^N, \theta_A)$, $u'_A = u_A(s', \theta_A)$, $u'_B = u_B(s', \theta_B)$, $s_B^O = s_B^O(s'_A, \theta_A)$ and $u_B^O = u_B(s'_A, \theta_B)$. Player B offers action s'_A and a monetary value of $\gamma \delta_B(s'_A) = \gamma(u'_B - u_B^O)$ to player A . Two cases can occur.

Player A accepts: $u'_A + \gamma \delta_B(s'_A) \geq u_A^N$. Strategy (s'_A, s'_B) is played.

$$\begin{aligned} PoA^A &\leq \frac{u_A^N + u'_B}{u'_A + u'_B} \leq \frac{u'_A + u'_B + \gamma \cdot u'_B}{u'_A + u'_B} \\ &= 1 + \gamma \frac{u'_B}{u'_A + u'_B} \leq 1 + \gamma \end{aligned}$$

Player A rejects: $u'_A + \gamma \delta_B(s'_A) < u_A^N$. Strategy (s_A^N, s_B^O) is played.

$$PoA^R \leq \frac{u_A^N + u'_B}{u_A^N + u_B^O} \leq 1 + \frac{u'_B}{u_A^N + u_B^O} \leq 1 + \frac{1}{\gamma}$$

where the last inequality comes from:

$$u'_A + \gamma \delta_B(s'_A) < u_A^N \Leftrightarrow \gamma u'_B < u_A^N + \gamma u_B^O - u'_A < u_A^N + u_B^O$$

□

When $\gamma = 1$, the price of anarchy is two, but player B has no incentive to choose such a value. If $\gamma = 0.5$, the price of anarchy is three. Of course, player B will choose $\gamma' = \gamma^*(s'_A, \theta_B)$. Lemma 4 indicates that the worst-case outcome is $(1 + \gamma')$ when player A accepts with a probability $P(s'_A, \gamma', \theta_B)$ and $(1 + \frac{1}{\gamma'})$ otherwise. This yields the following result.

Theorem 2. The Bayesian Nash price of anarchy of the single-offer mechanism for one-way games is at most:

$$\frac{\gamma' + 1}{\gamma'} (1 - P(s'_A, \gamma', \theta_B)(1 - \gamma'))$$

where:

$$\gamma' = \arg\text{-max}_{\gamma \in \mathbb{R}_{[0,1]}} P(s'_A, \gamma, \theta_B)(1 - \gamma)$$

Proof. By combining Lemmas 3 and 4, we can derive the following upper bound for the PoA.

$$\begin{aligned}
 PoA &\leq P(s'_A, \gamma', \theta_B) PoA^A(\gamma') + (1 - P(s'_A, \gamma', \theta_B)) PoA^R(\gamma') \\
 &= P(s'_A, \gamma', \theta_B) (1 + \gamma') + (1 - P(s'_A, \gamma', \theta_B)) \left(1 + \frac{1}{\gamma'}\right) \\
 &= 1 + \frac{1}{\gamma'} + P(s'_A, \gamma', \theta_B) \left(\gamma' - \frac{1}{\gamma'}\right) \\
 &= \frac{\gamma' + 1}{\gamma'} + P(s'_A, \gamma', \theta_B) \left(\frac{\gamma'^2 - 1}{\gamma'}\right) \\
 &= \frac{\gamma' + 1}{\gamma'} (1 - P(s'_A, \gamma', \theta_B) (1 - \gamma'))
 \end{aligned}$$

□

To get a better idea of how the mechanism improves the social welfare, it is useful to quantify the price of anarchy in Theorem 2 for a specific class of distributions.

Corollary 1. If $\delta_A(s'_A, \theta_A)$ has a cumulative distribution function $F(x) = (x/\delta_B)^\beta$ between zero and δ_B , with $0 < \beta \leq 1$, then $\gamma = \frac{\beta}{\beta+1}$, and the price of anarchy is at most:

$$\left(2 + \frac{1}{\beta}\right) (1 - \beta^\beta (1 + \beta)^{-(\beta+1)})$$

For example, if $\beta = 1$, then $F(x)$ is the uniform distribution, $\gamma = \frac{1}{2}$, and the expected price of anarchy is at most 2.25.

This corollary, in conjunction with Lemma 1, gives us the cost of enforcing individual rationality, moving from a price of anarchy of two to a price of 2.25 in the case of a uniform distribution.

The strategy $\langle s'_A, \gamma' \rangle$ is of independent interest. It indicates how a player with limited computational power can achieve an outcome that satisfies individual rationality without optimizing overall strategies.

6. Multi-Offer Mechanism

This section extends the single-offer mechanism by allowing player B to make multiple monetary offers for the same proposed action. Our main result shows that making counteroffers under commitment does not improve efficiency over the single-offer mechanism. By commitment, we mean that player B must be able to guarantee that the price schedule she originally announces will not be modified in the future. In this setting, incentive-compatibility and individual rationality conditions refer to the optimality of each player’s complete plan of action, in which players initially commit directly to entire strategies.

The single-offer mechanism was characterized by an action $s_A \in \mathcal{S}_A$ and a single value $\gamma \in \mathbb{R}_{[0,1]}$. The multi-offer mechanism is characterized by a four-tuple:

$$(s_A \in \mathcal{S}_A, n, \gamma = (\gamma_1, \dots, \gamma_n) \in \mathbb{R}_{[0,1]}^n, p = (p_1, \dots, p_n) \in \mathbb{R}_{[0,1]}^n)$$

where n is the number of offers, $(\gamma_1, \dots, \gamma_n)$ is a sequence of numbers in $\mathbb{R}_{[0,1]}^n$ to compute the ratios of $\delta_B(s_A, \theta_B) = u_B(s_B(s_A, \theta_B), \theta_B) - u_B^O(s_A, \theta_B)$ to be offered and (p_1, \dots, p_n) is a sequence of probabilities for continuing to make offers where we assume that $p_1 = 1$. The multi-offer mechanism is defined as follows:

1. Player B selects an action $s_A \in \mathcal{S}_A$ to propose to player A .
2. Player B also computes her outside option $s_B^O(s_A, \theta_B)$ in case player A rejects action s_A , and we denote by $u_B^O(s_A, \theta_B)$ the expected payoff from her outside option.
3. Player B selects $\gamma = (\gamma_1, \dots, \gamma_n) \in \mathbb{R}_{[0,1]}^n$ and $p = (p_1, \dots, p_n) \in \mathbb{R}_{[0,1]}^n$, with $p_1 = 1$. Player B has to commit to this sequence of values (despite what she learns from player A 's actions).
4. At step $1 \leq i \leq n$, player B proposes a monetary value of $\gamma_i \cdot \delta_B(s_A, \theta_B)$ with $\delta_B(s_A, \theta_B) = u_B(s_B(s_A, \theta_B), \theta_B) - u_B^O(s_A, \theta_B)$ and $\gamma_i \in \mathbb{R}_{[0,1]}$ to player A in the hope that she accepts to play strategy s_A instead of strategy s_A^N .
5. Player A decides whether to accept the offer.
6. If player A accepts the offer, the outcome of the game is $(s_A, s_B(s_A, \theta_B))$.
7. If player A rejects the offer, set $i \leftarrow i + 1$, and go to Step 4 with probability p_i .
8. Otherwise, the outcome of the game is the outside option $(s_A^N(\theta_A), s_B^O(s_A, \theta_B))$.

For ease of notation, we denote $\delta_B(s_A) = \delta_B(s_A, \theta_B)$ and $\delta_A(s_A) = \delta_A(s_A, \theta_A)$ for the rest of this section, where $\delta_A(s_A, \theta_A) = u_A(s_A^N(\theta_A), \theta_A) - u_A(s_A, \theta_A)$. In the multiple-offer mechanism, player B makes a sequence of offers $\gamma_i \delta_B(s_A)$ to player A to play strategy s_A . The first offer is $\gamma_1 \delta_B(s_A)$. If player A refuses the offer, then player B makes a second offer $\gamma_2 \delta_B(s_A)$ with probability p_2 . Hence, with probability $1 - p_2$, player B makes no offer, and the outcome of the game is $(s_A^N(\theta_A), s_B^O(s_A, \theta_B))$. In general, at iteration i , player B makes an offer $\gamma_i \delta_B(s_A)$ with probability p_i and the outside option is played with probability $1 - p_i$. The mechanism stops when player A accepts an offer or when player B stops making offers to player A . In this last case, once again, the outside option is played.

Observe that player A could reject an offer even if it is more profitable than playing her maximizing utility action $s_A^N(\theta_A)$ because she may expect a better offer in the future. To avoid this behavior, the multi-offer mechanism imposes a condition on the γ_i 's and p_i 's to ensure that player A accepts the first offer that gives her a higher payoff than her default action $s_A^N(\theta_A)$. Two conditions must hold for player A to accept an offer in step $i \in [1, \dots, n]$:

(a) Individual rationality:

$$\gamma_i \delta_B(s_A) \geq \delta_A(s_A) \tag{3}$$

which is equivalent to Proposition 2.

(b) Greater expected utility in step i than in step $i + 1$:

$$\gamma_i \delta_B(s_A) + u_A(s_A, \theta_A) \geq p_{i+1} (\gamma_{i+1} \delta_B(s_A) + u_A(s_A, \theta_A)) + (1 - p_{i+1}) u_A(s_A^N(\theta_A), \theta_A)$$

which is equivalent to:

$$\frac{\gamma_i - p_{i+1} \gamma_{i+1}}{1 - p_{i+1}} \delta_B(s_A) \geq \delta_A(s_A) \tag{4}$$

We now show that the multiple-offer mechanism is in fact equivalent to the single-offer mechanism. We use the notation:

$$S_i = \begin{cases} 0 & i = 0 \\ \frac{\gamma_i - p_{i+1} \gamma_{i+1}}{1 - p_{i+1}} & n > i > 0 \\ \gamma_n & i = n \end{cases}$$

so that Condition (4) can be expressed as $S_i \delta_B(s_A) \geq \delta_A(s_A)$.

Note that if player A refuses an offer with γ_i , she will also refuse offers with smaller ratios. This observation leads to the following proposition.

Proposition 4. *In the multi-offer mechanism,*

$$\gamma_{i+1} > \gamma_i, \forall i \in [1, \dots, n - 1]$$

Therefore, Proposition 4 states that counteroffers should be increasing with time.

Lemma 5. *In the multiple-offer mechanism, for all $i \in [1, \dots, n]$,*

$$\gamma_i \geq S_i$$

Proof. Assume that $\gamma_i < S_i$. By the definition of S_i , it follows that $\gamma_i - p_{i+1}\gamma_i < \gamma_i - p_{i+1}\gamma_{i+1}$, and hence, $\gamma_i > \gamma_{i+1}$. This contradicts Proposition 4, stating that the γ 's are defined as a non-decreasing sequence. \square

Corollary 2. *Condition (4) implies Condition (3).*

Proof. $S_i \leq \gamma_i$ and $\delta_A(s_A) \leq S_i \delta_B(s_A)$ implies $\delta_A(s_A) \leq S_i \delta_B(s_A) \leq \gamma_i \delta_B(s_A)$. \square

If player A rejected the offer in step $i - 1$, then Conditions (3) and (4) were both not satisfied in step $i - 1$. By Lemma 5, only two cases may occur:

1. If $\gamma_i \delta_B(s_A) < \delta_A(s_A)$, then it must be the case that:

$$S_{i-1} \delta_B(s_A) \leq \gamma_{i-1} \delta_B(s_A) < \delta_A(s_A) \tag{5}$$

2. If $\gamma_i \delta_B(s_A) \geq \delta_A(s_A)$, then:

$$S_{i-1} \delta_B(s_A) < \delta_A(s_A) \leq \gamma_{i-1} \delta_B(s_A) \tag{6}$$

The disjunction of Conditions (5) and (6) yields the following inequality:

$$S_{i-1} \delta_B(s_A) < \delta_A(s_A) \tag{7}$$

By Corollary 2, if player A accepts in step i given that she rejected in step $i - 1$, we have that:

$$\delta_A(s_A) \leq S_i \delta_B(s_A) \tag{8}$$

Recalling Definition 8, the cumulative distribution function of random variable $\delta_A(s_A)$ was denoted by,

$$P(s_A, \gamma) = \Pr [\delta_A(s_A) \leq \gamma \delta_B(s_A)]$$

Hence, the probability of acceptance in step i can be derived from Conditions (7) and (8) as:

$$\Pr[S_{i-1} \delta_B(s_A) < \delta_A(s_A) \leq S_i \delta_B(s_A)] = P(s_A, S_i) - P(s_A, S_{i-1})$$

Player B aims at choosing the γ_i 's, the probabilities p_i 's and action s_A to maximize her expected utility, which is equivalent to the following optimization problem.

$$\max_{\gamma, p} \sum_{i=1}^n \left[\left(\prod_{j=1}^i p_j \right) (P(s_A, S_i) - P(s_A, S_{i-1})) (1 - \gamma_i) \right] \tag{9}$$

$$\text{s.t. } p_1 = 1 \tag{10}$$

$$S_0 = 0 \tag{11}$$

$$S_i = \frac{\gamma_i - p_{i+1}\gamma_{i+1}}{1 - p_{i+1}} \tag{12}$$

$$S_n = \gamma_n \tag{13}$$

$$S_1 \leq S_2 \leq \dots \leq S_n \leq 1 \tag{14}$$

where the term $\prod_{j=1}^i p_j$ is the probability of reaching the i -th offer. We are now ready to state the main result of this section.

Theorem 3. *The multi-offer mechanism is equivalent to the single-offer mechanism in one-way games.*

Proof. By Equation (12),

$$(1 - p_{i+1})(1 - S_i) = (1 - \gamma_i) - p_{i+1}(1 - \gamma_{i+1})$$

Then, by using Equation (13) and grouping the $P(s_A, S_i)$ terms, the objective function becomes:

$$\sum_{i=1}^{n-1} \left[\left(\prod_{j=1}^i p_j \right) (1 - p_{i+1}) P(s_A, S_i) (1 - S_i) \right] + \left(\prod_{j=1}^n p_j \right) P(s_A, \gamma_n) (1 - \gamma_n) \tag{15}$$

Observe that each term in the objective function features an expression of the form $P(s_A, x)(1 - x)$. Hence, the objective is bounded by above by:

$$\text{Equation(15)} \leq \sum_{i=1}^{n-1} \left[\left(\prod_{j=1}^i p_j \right) (1 - p_{i+1}) \cdot C \right] + C \cdot \prod_{j=1}^n p_j$$

where $C = \max_x P(s_A, x)(1 - x)$. We show that, for any given probabilities p , there is a unique solution that meets this upper bound. Let $x^* = \arg \max_x P(s_A, x)$. The right-hand term in Equation (15) is optimized by setting $\gamma_n = x^*$. We show by induction that all of the other terms are optimized by setting $\gamma_i = x^*$. Assume that this holds for $\gamma_{i+1}, \dots, \gamma_n$. We need to optimize $P(s_A, S_i)(1 - S_i)$. By induction,

$$S_i = \frac{\gamma_i - p_{i+1}x^*}{1 - p_{i+1}}$$

and assigning x^* to γ_i gives $S_i = x^*$ and $P(s_A, S_i)(1 - S_i) = C$. Since all γ_i are equal, this concludes the proof. \square

The above derivation is related to a well-known result from Sobel and Takahashi [13], which models an iterative bargaining where there is a buyer with a private reservation price and a seller with reservation price zero who makes all of the offers. There is a known fixed discount factor for each player, and

when these discount factors are equal (this is equivalent to have a probability for the next offer), they showed that, under commitment, the infinite horizon bargaining is equivalent to the single shot. There are differences between their model and ours: in our model, the buyer is making the offers; the probabilities are not fixed *a priori* (player *B* can choose them); and both outside options are private.

7. Conclusions

In one-way games, the utility of one player does not depend on the decisions of the other player. We showed that, in this setting, the outcome of a Nash equilibrium when no side payments are allowed can be arbitrarily far from the social welfare solution. When it is far enough, player *B* has bargaining power to incentivize player *A* monetarily, so that she moves from her equilibrium and cooperates to overcome bad social welfare. We have explored this possibility by analyzing the social welfare when side payments are allowed. In the setting with private information and when side payments are allowed, we proved that it is impossible to design a Bayes–Nash incentive-compatible mechanism for one-way games that is budget-balanced, individually rational and efficient. To alleviate these negative results, we proposed two privacy-preserving mechanisms, a single-offer and a multi-offer mechanism, and showed that both are equivalent in terms of the equilibrium outcomes reached.

The single-offer mechanism is simple for both parties, as well as for the broker who just makes sure that the players follow the protocol. This mechanism also requires minimal information from the agents who perform all of the combinatorial computations, while it incentivizes them to cooperate towards social welfare in a distributed setting. Moreover, the mechanism has the following desirable properties: it is budget-balanced and satisfies the individual rationality constraints and Bayesian incentive-compatibility conditions. Whereas the mechanism can only guarantee interim individual rationality for player *B* since her outside option is not known a priori, it provides *ex post* individual rationality for player *A*. Additionally, we showed that, in a realistic setting, where agents have limited computational resources, a simpler version of the mechanism can be implemented without overly deteriorating the social welfare.

It is an open question whether there exists another mechanism (possibly more complex) that could lead to a better efficiency, while keeping the above properties. Indeed, in one-way games, player *A* has an intrinsic advantage over player *B*, which is not easy to overcome. One possible promising mechanism consists of player *B* setting rewards for all player *A*'s actions and player *A* choosing one in return for that money. This is known as the Bayesian unit-demand item-pricing problem (BUPP) [14]. Recent work has shown this problem to be NP-hard [15], but a factor three approximation to the optimal expected revenue of player *B* is obtained in [14] (subsequently improved to two in [16]). In the context of our paper, several interesting questions arise from the Bayesian unit-demand item-pricing problem: What is the efficiency achieved by the BUPP in one-way games? What is the impact of a constant factor approximation for the revenue on the social welfare in one-way games?

Acknowledgments

National Information and Communications Technology Australia (NICTA) is funded by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre of Excellence Program.

Author Contributions

All authors contributed equally to this article.

Conflicts of Interest

The authors declare no conflict of interest.

References

1. Abeliuk, A.; Berbeglia, G.; Van Hentenryck, P. A Bargaining Mechanism for One-Way Games. In Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI), Buenos Aires, Argentina, 25–31 July 2015.
2. Coase, R.H. The problem of social cost. *J. Law Econ.* **1960**, *3*, 1–44.
3. Hahnel, R.; Sheeran, K.A. Misinterpreting the Coase theorem. *J. Econ. Issues* **2009**, *43*, 215–238.
4. Gentle, J.E. *Computational Statistics*; Springer: Berlin, Germany, 2009; Volume 308.
5. Jackson, M.O.; Wilkie, S. Endogenous games and mechanisms: Side payments among players. *Rev. Econ. Stud.* **2005**, *72*, 543–566.
6. Myerson, R.B.; Satterthwaite, M.A. Efficient mechanisms for bilateral trading. *J. Econ. Theory* **1983**, *29*, 265–281.
7. D'Aspremont, C.; Gérard-Varet, L.A. Incentives and incomplete information. *J. Public Econ.* **1979**, *11*, 25–45.
8. Arrow, K. The Property Rights Doctrine and Demand Revelation Under Incomplete Information. In *Economics and Human Welfare: Essays in Honour of Tibor Scitovsky*; Boskin, M., Ed.; Academic Press: New York, NY, USA, 1979.
9. Williams, S.R. A characterization of efficient, Bayesian incentive compatible mechanisms. *Econ. Theory* **1999**, *14*, 155–180.
10. Krishna, V.; Perry, M. *Efficient Mechanism Design*; Working Paper; Department of Economics, The Pennsylvania State University: University Park, PA, USA, 1998.
11. Kennan, J.; Wilson, R. Bargaining with private information. *J. Econ. Lit.* **1993**, *31*, 45–104.
12. Chatterjee, K.; Samuelson, W. Bargaining under incomplete information. *Oper. Res.* **1983**, *31*, 835–851.
13. Sobel, J.; Takahashi, I. A multistage model of bargaining. *Rev. Econ. Stud.* **1983**, *50*, 411–426.
14. Chawla, S.; Hartline, J.D.; Kleinberg, R. Algorithmic pricing via virtual valuations. In Proceedings of the 8th ACM Conference on Electronic Commerce (ACM), San Diego, CA, USA, 11–15 June 2007; pp. 243–251.

15. Chen, X.; Diakonikolas, I.; Paparas, D.; Sun, X.; Yannakakis, M. The Complexity of Optimal Multidimensional Pricing. In Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SIAM), SODA '14, Portland, OR, USA, 5–7 January 2014; pp. 1319–1328.
16. Chawla, S.; Hartline, J.D.; Malec, D.L.; Sivan, B. Multi-parameter Mechanism Design and Sequential Posted Pricing. In Proceedings of the 42nd ACM Symposium on Theory of Computing (ACM), STOC '10, Cambridge, MA, USA, 6–8 June 2010; pp. 311–320.

© 2015 by the authors; licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution license (<http://creativecommons.org/licenses/by/4.0/>).