Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015)

A Bargaining Mechanism for One-Way Games

Andrés Abeliuk
NICTA and
University of Melbourne
andres.abeliuk@nicta.com.au

Gerardo Berbeglia
NICTA and
Melbourne Business School
g.berbeglia@mbs.edu

Pascal Van Hentenryck
NICTA and
Australian National University
pvh@nicta.com.au

Abstract

We introduce one-way games, a framework motivated by applications in large-scale power restoration, humanitarian logistics, and integrated supply-chains. The distinguishable feature of the games is that the payoff of some player is determined only by her own strategy and does not depend on actions taken by other players. We show that the equilibrium outcome in one-way games without 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. Finally, we propose a privacy-preserving mechanism that is incentive-compatible and budget-balanced, satisfies ex-post individual rationality conditions, and produces an outcome which is more efficient than the equilibrium without payments.

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 [Coase, 1960] 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 (2009) 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.

Other real-life applications are found in large-scale restoration of interdependent infrastructures after significant disruptions [Cavdaroglu *et al.*, 2013; Coffrin *et al.*, 2012], humanitarian logistics over multiple states or regions [Van Hentenryck *et al.*, 2010], supply chain coordination (see, e.g., Voigt (2011)), integrated logistics, and the joint planning and the control of gas and electricity networks. Consider, for instance, the restoration of the power system and the telecommunication network after a major disaster. As explained in [Cavdaroglu *et al.*, 2013], there are one-way dependencies between the power system and the telecommunication network. This means, for instance, that some power lines must be restored before some parts of the telecommunication network become available. It is possible to use centralized mechanisms for restoring the system as a whole. However, it is often the case that these restorations are performed by different agencies with independent objectives and selfish behavior may have a strong impact on the social welfare. It is thus important to study whether it is possible to find high-quality outcomes in decentralized settings when stakeholders proceed independently and do not share complete information about their utilities.

This paper aims at taking a first step in this direction by proposing a class of two players one-way dependent decision settings which abstracts some of the salient features of these applications and 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 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, incentive-compatible, and have a small price of anarchy. Such mechanisms are particularly useful for rational agents who have incomplete information about other agents. Our main positive result is a single-offer bargaining mechanism which under reasonable assumptions on the players, increases the social welfare compared to the setting where no side payments are allowed.

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 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 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.,

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

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

One-way games assumes that players simultaneously choose their actions and players are risk-neutral agents. Hence, if side payments are not allowed, player A will play an action s_A^N that yields her a maximum payoff, i.e.,

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

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

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

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 its optimal payoff, but our motivating applications aim at optimizing a global welfare function

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

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 θ is defined as*

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

Throughout this paper, we use the following simple running example to illustrate key concepts.

Example 1. Consider the instance where player A has two possible actions $s_A^1, s_A^2 \in \mathcal{S}_A$. Action s_A^1 has a payoff distributed according to a uniform distribution between 0 and 100, while action s_A^2 has a constant payoff of 100. The set of dominant actions for player B corresponds to the set of best responses, i.e., $s_B(s_A^1)$ and $s_B(s_A^2)$. We set payoffs to be $u_B(s_B(s_A^1)) = x$ and $u_B(s_B(s_A^2)) = 0$, 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 $50 + x$. Thus, the price of anarchy is $\frac{50+x}{100}$ if $x \geq 50$ and 1 otherwise. Notice that the PoA is unbounded on x for the Nash equilibrium.

We now quantify the price of anarchy in one-way games.

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

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

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

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

□

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

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 (2005) 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. *One of the key insights of this paper is the recognition that this result does not hold under incomplete information.* To remedy this negative result, we present a bargaining mechanism.

3 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 which 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 2. *A social choice function is ex-post efficient if, for all $\theta \in \Theta$, we have $k(\theta) \in \text{arg-max}_{s \in \mathcal{S}} \sum_i u_i(s, \theta)$.*

Definition 3. *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 4. *A social choice function is Bayes-Nash incentive compatible (IC) if for every player i:*

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

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

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

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

In the context of one-way games, both players have positive outside options that depend only in 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 $\overline{u}_A(\theta_A) = u_A(s_A^N(\theta_A), \theta_A)$ and $\overline{u}_B(\theta_B) = u_B(s^N(\theta), \theta_B)$ respectively.

3.1 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 (1983) 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 gives every agent type non-negative expected gains from participation (i.e., ex interim individual rationality).

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

Consider the Myerson-Satterthwaite bilateral bargaining setting.

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 (resp. 2) has a probability density distribution $f_2(v_2)$ (resp. $f_1(v_1)$) 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 $\sigma : V_1 \times V_2 \to [0, 1]$ that specifies the probability that the object is transferred from the seller to the buyer and (2) a monetary transfer scheme $p : V_1 \times V_2 \to \mathbb{R}^2$. In this setting, ex-post efficiency is achieved if $\sigma(v_1, v_2) = 1$ when $v_1 < v_2$, and 0 otherwise.

Our result consists in 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 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 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 Myerson-Satterthwaite impossibility result, which concludes our proof. □

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, VCG and dAGVA [d'Aspremont and Gérard-Varet, 1979; Arrow, 1979] 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 (1999) and Krishna and Perry (1998) for alternative derivations of the impossibility result for bilateral trading under the Groves family of mechanisms.

4 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 (1993).

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 leads to the following lemma.

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

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

In the one-way game, strategy $s'(\theta)$ has a price of anarchy of 2 (i.e., $\forall \theta \; PoA(\theta) = 2$).

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 $SW(s', \theta) < SW(s^N, \theta)$, 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 (1983). 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 provide quality guarantee on the mechanism outcome. Our single-offer mechanism is defined as follows:

Single-offer mechanism:

1. Player B selects an action $s_A \in \mathcal{S}_A$ to propose to player A.
2. Player B also computes her expected outside option $u_B^N(\theta_B) = \mathbb{E}_{\theta_A|\theta_B}\left[u_B(s^N(\theta), \theta_B)\right]$.
3. Player B proposes a monetary value of $\gamma \cdot \Delta_B(s_A, \theta_B)$ with $\Delta_B(s_A, \theta_B) = u_B(s_B(s_A, \theta_B), \theta_B) - u_B^N(\theta_B)$ 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 Nash equilibrium $\left(s_A^N(\theta_A), s_B^N(\theta_B)\right)$.

It is worth observing that a broker is required in this mechanism to ensure that the outcome $\left(s_A^N(\theta_A), s_B^N(\theta_B)\right)$ 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).

Proposition 2. *If players A and B play the single-offer mechanism, player A accepts the offer whenever*

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

We have designed a mechanism that satisfies individual rationality: Player B never offers more than $\Delta_B(s_A, \theta_B)$ and her payoff is never worse than her expected outside option. By Proposition 2, player A is always better off playing the single-offer mechanism.

Example 2. *(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 accepts the offer if $c + u_A(s_A^1) \geq u_A(s_A^2) = 100$. Given that $u_A(s_A^1)$ follows a uniform distribution, the probability that player A accepts the offer is $\frac{c}{100}$ and such an offer has an expected payoff of $\frac{c}{100} \cdot (x - c)$ for player B. The optimal value for the offer is given by $c^* = \frac{x}{2}$. This leads to an expected social welfare $SW = \frac{c^*}{100} \cdot (50 + x) + (1 - \frac{c^*}{100}) \cdot 100$ for the single-offer mechanism. Recall that the optimal social welfare is $50 + x$ if $x \geq 50$ and 100 otherwise. Therefore, the mechanism has a price of anarchy of $\frac{50+x}{SW} \leq 1.21$ if $x \geq 50$ and $\frac{100}{SW} \leq 1.04$ otherwise. This contrasts with the unbounded PoA obtained by the Nash equilibrium when no side payments are allowed.

We now generalize the analysis done in Example 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 $\Delta_A(s_A, \theta_A) \leq \gamma \Delta_B(s_A, \theta_B)$, 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 $u_B^N(\theta_B)$.

Definition 6. *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*

$$\mathbb{E}_{\theta_A}\left[U_B(s_A, \gamma, \theta_B)\right] = u_B^N(\theta_B) + P(s_A, \gamma, \theta_B)\left((1 - \gamma) \cdot \Delta_B(s_A, \theta_B)\right),$$

$$\mathbb{E}_{\theta_A}\left[U_A(s_A, \gamma, \theta_B)\right] = \mathbb{E}_{\theta_A}[u_A^N(\theta_A)] + P(s_A, \gamma, \theta_B) \cdot \left(\gamma \Delta_B(s_A, \theta_B) - \mathbb{E}_{\theta_A}\left[\Delta_A(s_A, \theta_A)\right]\right),$$

where the probability that player A accepts to play s_A is defined by

$$P(s_A, \gamma, \theta_B) = \Pr\left[\gamma \cdot \Delta_B(s_A, \theta_B) \geq \Delta_A(s_A, \theta_A)\right].$$

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*

$$s_A^*(\theta_B) = \underset{s_A}{\text{arg-max}}\ P\left(s_A, \gamma^*(s_A, \theta_B), \theta_B\right) \cdot \left(1 - \gamma^*(s_A, \theta_B)\right) \cdot \Delta_B(s_A)$$

where

$$\gamma^*(s_A, \theta_B) = \underset{\gamma}{\text{arg-max}}\ P(s_A, \gamma, \theta_B) \cdot (1 - \gamma).$$

4.1 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 $s'_A = \text{arg-max}_{s_A \in \mathcal{S}_A} u_B(s_B(s_A, \theta_B), \theta_B)$. 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. *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

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

Two cases can occur. The first case is

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

i.e., the probabilty 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.,

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

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

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

Similarly, consider γ^{**} such that

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

which implies

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

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 with strategy $\langle s'_A, \gamma' \rangle$. □

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 case of agreement and disagreement of player A.

Lemma 4. *Consider action $s' = \arg\max_{s \in \mathcal{S}} u_B(s, \theta_B)$ 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,*

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

Proof. Define $u^N_A = u_A(s^N_A, \theta_A)$, $u'_A = u_A(s', \theta_A)$ and $u'_B = u_B(s', \theta_B)$. Two cases can occur.

Case $u'_A + \gamma\Delta_B(s'_A) \geq u^N_A$. Strategy (s'_A, s'_B) is played.

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

Case $u'_A + \gamma\Delta_B(s'_A) < u^N_A$. Player A plays s^N_A.
$\gamma \cdot u^N_B \leq u'_A + \gamma \cdot u^N_B < u^*_A + \gamma \cdot \underline{u}_B \leq u^*_A + \underline{u}_B$ then,

$$\begin{aligned} PoA^R &\leq \frac{u^N_A + u'_B}{u^N_A + u^N_B} \leq 1 + \frac{u'_B}{u^N_A + u^N_B} \\ &\leq 1 + \frac{u'_B}{\gamma \cdot u'_B} = 1 + \frac{1}{\gamma}. \end{aligned}$$

□

When $\gamma = 1$, the price of anarchy is 2 but player B has no incentive to choose such a value. If $\gamma = 0.5$, the price of anarchy is 3. 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 price of anarchy for one-way games is at most*

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

where

$$\gamma' = \operatorname*{arg\text{-}max}_{\gamma} P(s'_A, \gamma, \theta_B)(1 - \gamma).$$

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 0 and Δ_B, with $0 < \beta \leq 1$, then $\gamma = \frac{\beta}{\beta+1}$ and the price of anarchy is at most*

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

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 2 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 over all strategies.

5 Conclusion

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 can be arbitrarily far from the social welfare solution. We also 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 a privacy-preserving mechanism based on a single-offer.

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 the combinatorial computations, while it incentivizes them to cooperate towards the 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. 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 a intrinsic advantage over player B, which is not easy to overcome. There are also many other directions for future research. It is important to generalize one-way games to multiple players. Moreover, there are applications where the dependencies are in both directions, e.g., the restoration of the power and the gas systems considered in Coffrin et al. (2012). These applications typically have multiple components to restore and the dependencies form an acyclic graph. Hence such a mechanism would likely need to consider this internal structure to obtain efficient outcomes.

Acknowledgments

NICTA is funded by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre of Excellence Program.

References

[Arrow, 1979] K.J. Arrow. The property rights doctrine and demand revelation under incomplete information. *In Boskin, M. (Ed.), Economics and Human Welfare*, 1979.

[Cavdaroglu *et al.*, 2013] Burak Cavdaroglu, Erik Hammel, John E Mitchell, Thomas C Sharkey, and William A Wallace. Integrating restoration and scheduling decisions for disrupted interdependent infrastructure systems. *Annals of Operations Research*, 203:1–16, 2013.

[Chatterjee and Samuelson, 1983] Kalyan Chatterjee and William Samuelson. Bargaining under incomplete information. *Operations Research*, 31:835–851, 1983.

[Coase, 1960] Ronald Harry Coase. The problem of social cost. *The Journal of Law and Economics*, 3:1–44, 1960.

[Coffrin *et al.*, 2012] Carleton Coffrin, Pascal Van Hentenryck, and Russell Bent. Last-mile restoration for multiple interdependent infrastructures. In *In Proc. 26th AAAI Conference on Artificial Intelligence (AAAI)*, pages 455–463, 2012.

[d'Aspremont and Gérard-Varet, 1979] Claude d'Aspremont and Louis Andre Gérard-Varet. Incentives and incomplete information. *Journal of Public Economics*, 11:25–45, 1979.

[Hahnel and Sheeran, 2009] Robin Hahnel and Kristen A Sheeran. Misinterpreting the coase theorem. *Journal of Economic Issues*, 43:215–238, 2009.

[Jackson and Wilkie, 2005] Matthew O Jackson and Simon Wilkie. Endogenous games and mechanisms: Side payments among players. *The Review of Economic Studies*, 72:543–566, 2005.

[Kennan and Wilson, 1993] John Kennan and Robert Wilson. Bargaining with private information. *Journal of Economic Literature*, 31:45–104, 1993.

[Krishna and Perry, 1998] Vijay Krishna and Motty Perry. Efficient mechanism design. *Available at SSRN 64934*, 1998.

[Myerson and Satterthwaite, 1983] Roger B Myerson and Mark A Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29:265–281, 1983.

[Van Hentenryck *et al.*, 2010] Pascal Van Hentenryck, Russell Bent, and Carleton Coffrin. Strategic planning for disaster recovery with stochastic last mile distribution. In *CPAIOR*, pages 318–333. Springer, 2010.

[Voigt, 2011] Guido Voigt. *Supply Chain Coordination in Case of Asymmetric Information*. Springer, 2011.

[Williams, 1999] Steven R Williams. A characterization of efficient, bayesian incentive compatible mechanisms. *Economic Theory*, 14(1):155–180, 1999.