Currency Systems for Efficient and Equitable Exchange Among Selfish Agents in the Air Traffic Management Domain

Geert Jonker,
Frank Dignum,
John-Jules Ch. Meyer

Department of Information and Computing Sciences, Utrecht University

Technical Report UU-CS-2008-004

www.cs.uu.nl

ISSN: 0924-3275

Currency Systems for Efficient and Equitable Exchange Among Selfish Agents in the Air Traffic Management Domain *

Geert Jonker, Frank Dignum, John-Jules Ch. Meyer
{geertj,dignum,jj}@cs.uu.nl
Dep. of Information and Computing Sciences
Utrecht University

February 14, 2008

Abstract

We study collaboration among selfish agents in the air traffic management domain, in particular in the tactical airport planning domain. Collaboration can be seen as a social exchange scenario, in which the efforts of performing tasks are the resources that are being exchanged. We investigate the use of money as a facilitator for distributed social exchange. We give conditions under which a market mechanism with the use of standard currency leads to efficient and equitable exchange among benevolent agents. We show that with selfish agents, under the same conditions efficiency and equitability are not guaranteed. We investigate whether a coalition of benevolent agents can penalize malicious agents such that efficiency and equitability are restored. We show that a straightforward penalty rule is not robust, as it is attractive to deviate from it.

In the second part of this paper, we present a novel currency system that facilitates efficient and equitable exchange among selfish agents. The mechanism allows each agent to issue its own money and circulate that of others. Each credit is valuated individually, based on the reputations of the agents that have used the credit. We give a condition on the reputations under which the system successfully penalizes malicious agents. Unlike the standard currency system, this system is robust as it is not attractive to deviate from the penalty rule.

1 Introduction

There are two important trends in current day *air traffic management* (ATM) research. Firstly, a lot of research is done on what is called *collaborative decision making* [3, 10]. The aim of this research is to increase the amount and quality of information that is exchanged among the several parties involved in ATM. This will lead to an improved ability to plan operations, both in the short and the long run, and an improved ability to adapt plannings to occurring deviations.

Secondly, market-based techniques are more and more being investigated and applied in several areas, including slot allocation [4, 13], slot exchange [22, 3] and route charging [9, 6].

*This research is supported by the Technology Foundation STW, applied science division of NWO and the technology programme of the Ministry of Economic Affairs. Project DIT5780: Distributed Model Based Diagnosis and Repair

The need for new, more efficient techniques is strong, as airports and airways worldwide are at or near their maximal capacity. Nevertheless, most innovative proposals never actually make it to implementation. In our view, this is a result of thinking only in terms of the classical aim of ATM, namely *efficiency*. In this paper we argue that a focus on *equity* is needed too to make collaborative techniques successful. Equity is beginning to receive attention in the ATM literature [23], but it has not been thoroughly researched yet.

The domain of our research is the phase of planning known as *tactical airport planning*. This phase of planning is concerned with the sequencing of arriving and departing aircraft and their scheduling on the gates, as well as with various ground handling services. A predetermined plan exists, but deviations that occur at the last moment may make it infeasible. In that case the plan needs to be *repaired*. This can involve delaying flights, changing gates or runways, putting passengers or luggage on other flights, rescheduling ground services, etc.

Plan repair is currently done by human air traffic controllers. This significantly limits the complexity of repairs that can be constructed. As a rule of thumb, planning conflicts are resolved in such a way that all parties other than the responsible party are affected as little as possible. For instance, if an incoming flight is delayed and as a result its gate reservation is delayed and a conflict with the next gate reservation occurs, the first reservation is taken off the gate and scheduled at the next available spot in the planning. Collaboration among airlines occurs very little, firstly because many airport procedures are not designed for collaboration, and secondly because airlines are not eager to help each other as they have no guarantee that provided help will be reciprocated.

Efficiency is traditionally the main aim of plan repair. Plans should be repaired in such a way that the total effort incurred by the parties involved is minimal. However, the rule of thumb described above shows that equity is implicitly required as well; controllers try to avoid airlines being affected by problems they did not cause.

We view tactical airport planning as a *social exchange scenario*, in which airlines exchange favours by helping to solve each others problems. We give definitions for efficiency and equity in the plan repair scenario and show that the latter is a criterium that should be measured over a longer period of time. We propose an agent-based plan repair mechanism in which airlines are able to jointly elect efficient repairs and in which a minimum level of equity is guaranteed. A key assumption is that agents are self-interested; they will try to maximize their own utility at the cost of others if possible. Nevertheless, they also care about their relative utility. An agent would not agree on a collaboration mechanism that gives another agent a much larger advantage than itself. It is well known that equity is an important factor in joint decision making [27, 7].

We also study the applicability of money as a facilitator for reciprocal exchange. We show that standard currency is a good facilitator of exchange if agents are benevolent. If agents are selfish however, the use of standard currency may lead to significant inequity. As a solution, we propose a novel monetary system that can be used by selfish agents to achieve efficient and equitable exchange.

2 Efficient and equitable exchange

We model the plan repair problem as a *repeated resource allocation* problem. Let $A = \{1, 2, \ldots, k\}$ be the set of agents, with typical elements $i, j \in A$. In each round r, one agent, the *problem owner* $o_r \in A$, has caused a planning conflict. Let $\mathbf{o} = \langle o_1, o_2, o_3, \ldots \rangle$ be the infinite *owner vector* denoting problem owner o_r for round r.

Each round, there are a number of ways in which a planning conflict can be resolved. These are called *repair candidates*. We assume that there are m candidates in each round, with typical notation $c \in \{1, \ldots, m\}$. A repair candidate allocates a single task to a single agent. Different repair candidates involve different tasks, which have different utilities to the agents performing them. Let $\mathfrak{u} : A \times A \times \mathbb{N} \times \mathbb{N} \to \mathbb{R}$ be the *utility profile* that, given two agents i and j, a round r and a candidate c, gives the utility that agent i incurs in aid of agent j in round r in candidate c. Throughout this paper we consider $\mathfrak{u}$ given, and will therefore omit it in most definitions. We use notation $u_{i,j,r,c}$ to denote $\mathfrak{u}(i,j,r,c)$. The candidates we will consider consist of only a single task[1], which we assume to have a negative utility, i.e., $\forall r, c\colon \exists i\colon (u_{i,o_r,r,c} < 0 \ \wedge \ \forall j\colon \ j \neq i \to u_{j,o_r,r,c} = 0)$. The agent performing the task in a candidate is called the *actor*, denoted $a_{r,c}$. Performed tasks are always in aid of the problem owner, so we have $\forall r, c, i, j\colon \ j \neq o_r \to u_{i,j,r,c} = 0$. Also, we assume that plan repair starts in round 1 and require $\forall i, j, c\colon \ u_{i,j,0,c} = 0$. Note that we model candidates only in terms of the utilities that are incurred by the agents executing them.

We use the following derived notation for groups of agents

$$u_{I,J,r,c} = \sum_{i \in I} \sum_{j \in J} u_{i,j,r,c}$$

$$u_{r,c} = u_{A,A,r,c}$$

We will sometimes omit braces for a set with a single element, e.g. as in $u_{i,A,r,c}$. We will sometimes use a tuple notation for candidates, e.g. $\langle 0, -10, 0 \rangle$ denotes a candidate assigning a task to agent 2 with utility -10. The set of agents A, the utility profile $\mathfrak{u}$ and the sequence of owners $\mathbf{o}$ comprise a *plan repair problem set*, denoted $\pi = \langle A, \mathfrak{u}, \mathbf{o} \rangle$.

In every round, there is at least one candidate that assigns a task to the problem owner, usually with a relatively low utility, i.e., a utility more negative than the others. This is called a *default candidate*. Thus, for a default candidate d in round r, it holds that $a_{r,d} = o_r$.

2.1 Efficient exchange

We say that a candidate c in round r is *efficient* iff

$$\forall c'\colon \ c' \neq c \to u_{r,c'} \leq u_{r,c} \tag{1}$$

Each round, one of the candidates has to be elected. Let $\mathbf{e} = \langle e_1, e_2, e_3, \ldots \rangle$ be an *allocation*, being an infinite sequence of elected candidates $e_1, e_2, e_3, \ldots$ for rounds $1, 2, 3, \ldots$. Let

$$U_{i,r}(\mathbf{e}) = \sum_{t=1}^{r} u_{i,A,t,e_t} \qquad U_{A,r}(\mathbf{e}) = \sum_{t=1}^{r} u_{A,A,t,e_t}$$

be the *cumulative utility* of agent i up to and including round r and the *cumulative efficiency* of allocation $\mathbf{e}$ up to and including round r respectively.

Definition 1 (Efficient allocation). *An allocation* $\mathbf{e}$ *is* efficient in round r *iff for any other allocation* $\mathbf{e}'$ *it holds that* $U_{A,r}(\mathbf{e}) \geq U_{A,r}(\mathbf{e}')$.

[1] We chose to use single-task candidates to keep the complexity of the theory to a minimum. Nevertheless, the notation we will use is for the largest part also suitable for multi-task candidates. See also section 8 for a discussion.

Under the assumption that candidate sets are independent from elected candidates, we have

Proposition 2.1. *Let* $\mathbf{e}$ *be an allocation. Then,* $\mathbf{e}$ *is efficient in round* $r \iff \forall t \leq r: e_t$ *is efficient.*

Proof. ($\Rightarrow$) Suppose that there is a round t in which e_t is not efficient. Then there exists c^* such that $u_{t,c} > u_{t,e_r}$. It is easy to see that for the allocation $\mathbf{e}' = \langle e_1, e_2, \ldots, e_{t-1}, c^*, e_{t+1}, \ldots, e_n \rangle$, it holds that $U_{A,r}(\mathbf{e}') > U_{A,r}(\mathbf{e})$. Thus, allocation $\mathbf{e}$ is not efficient in round r.

($\Leftarrow$) Suppose that $\mathbf{e}$ is not efficient in round r. Then, $\exists \mathbf{e}' : U_{A,r}(\mathbf{e}') > U_{A,r}(\mathbf{e})$. Then, there exists t^* such that $u_{t,e'_{t^*}} > u_{t^*,e_{t^*}}$. Then, e^*_t is not efficient. □

2.2 Equitable exchange

Besides being a resource allocation problem, the plan repair problem can also be seen as a social exchange scenario. An airline will be willing to help another airline if it knows that its actions will be reciprocated. An exchange of effort can result that is beneficial for each of the participants. However, such an exchange should be equitable in the eyes of the participants.

Equity in social exchange scenario's has been thoroughly researched in the social science literature. In most influential papers, equitable exchange is defined as satisfying an equation of the form:

$$\frac{O_x}{I_x} = \frac{O_y}{I_y} \tag{2}$$

where I_x and O_x denote the inputs and outputs respectively of agent x that are perceived as relevant in the exchange [11, 1, 24]. Inputs are usually derived from differences in social status, education, race, sex, etc [11]. Outputs are usually interpreted as *gains minus costs* [11].

Interpreting (2) for the ATM case, we first observe that I_i is equal for each agent i. The fact that all airlines are considered equal shows for instance in the *first come first serve* rule that is widely used in airport operations. Also, airlines do not hesitate to go to court if they feel treated unfairly, often with success [19].

To interpret O_x, we need to determine what the gains and costs of the agents in the exchange are. Two interpretations can be defended. First, the *marginal utility gain* of an agent could be taken to be its output, i.e., $U_i(\mathbf{e}) - U_i(\mathbf{e}')$ where $\mathbf{e}$ is the allocation as a result of the exchange, and $\mathbf{e}'$ is an allocation that would have occurred if no collaboration took place. The main argument for this interpretation is probably that the marginal utility gain is that what an agent itself perceives as its output, its nett gain as a result of participating in the exchange.

Alternatively, one could only consider the utilities of executing tasks, and thus ignore utility gains of agents 'receiving' tasks. Two arguments can be given to defend this interpretation. Firstly, considering only utilities of executing tasks is considered equitable in real life. Consider the following illustration. My neighbour and I have agreed on an exchange of tasks; I fix her computer and she bakes a cake for me. It might be the case that by fixing her computer, she is able to send her application form in time, giving her the job of her life. In that case her utility gain would be much higher than mine. Nevertheless, the exchange can be considered to be equitable, since the effort spent by both sides is equal.

Secondly, determining the marginal utility gain can be highly impractical. Consider the following scenario. An aircraft runs out of fuel and needs to land immediately. For this it needs the collaboration of the aircraft in front of it, which has to go out of the landing sequence so that the first aircraft can take its place. Suppose that if the second aircraft does not cooperate, the first aircraft could not land and would crash. If they would equalize marginal utility gains, they would have to determine the utility of crashing, which would be a very low value. This results in a very high potential utility gain for the first aircraft, which somehow would have to be equalized. It would be very hard for the first aircraft to reciprocate the second aircraft by anything else than saving it in a similar situation. Clearly, determining and transferring such utilities is highly unpractical. It is far more likely that both parties will agree on the first aircraft reciprocating the help of the second by an action of comparable effort. For the two reasons given, we choose to consider the utilities of performing tasks to be the outputs in the exchange. Let

$$U_{i,r}^{out}(\mathbf{e}) = \sum_{t=1}^{r} u_{i,A\backslash i,t,e_t}$$
$$U_{i,r}^{in}(\mathbf{e}) = \sum_{t=1}^{r} u_{A\backslash i,i,t,e_t}$$
$$U_{i,r}^{self}(\mathbf{e}) = \sum_{t=1}^{r} u_{i,i,t,e_t}$$

be the total utility agent i provided to others, the total utility others provided to i and the total utility i provided to itself respectively in allocation $\mathbf{e}$ up to round r.

Definition 2 (Equitable exchange). *Allocation* $\mathbf{e}$ *is equitable in round* r *iff* $\forall i: U_{i,r}^{out}(\mathbf{e}) = U_{i,r}^{in}(\mathbf{e})$.

It will not always be possible to find or achieve an equitable allocation. In that case, we strive for an allocation that is minimally inequitable. Inequity is measured using the *nett transferred utility*, defined as follows for an agent i in allocation $\mathbf{e}$ in round r.

$$U_{i,r}^{nett}(\mathbf{e}) = U_{i,r}^{out}(\mathbf{e}) - U_{i,r}^{in}(\mathbf{e})$$

As a measure for inequity, we use the egalitarian notion that inequity is equal to the nett transferred utility of the worst off agent.

Definition 3 (Inequity of exchange). *Inequity of allocation* $\mathbf{e}$ *is defined as*

$$U_r^-(\mathbf{e}) = \min\{\, U_{1,r}^{nett}(\mathbf{e}), U_{2,r}^{nett}(\mathbf{e}), \ldots, U_{k,r}^{nett}(\mathbf{e}) \,\}$$

Note that the utilities of actions an agent performs to solve its own conflicts do not affect the inequity of the exchange. When it is clear which allocation is referred to, we will use abbreviations $U_{i,r} = U_{i,r}(\mathbf{e})$, $U_{i,r}^{out} = U_{i.r}^{out}(\mathbf{e})$, etc.

3 Market mechanism

A *market mechanism* is a mechanism in which users transfer resources in exchange for money, where prices are determined by means of some clearing mechanism, for instance an auction. From another perspective, a market mechanism is a distributed administrative system for social exchange. Agents transfer resources in exchange for money. If an agent has provided many resources to others, it is entitled to receive many resources in return. As it has accumulated a corresponding sum of money, it is able to get these resources in return. Thus, the monetary balance of an agent represents, in an ideal market, the amount of utility an agent is entitled to or should provide from the viewpoint of equity. The fact that the market mechanism is a distributed system follows from the fact that no central authority is needed to keep track of who is entitled to what. Each agent itself possesses a certain amount of money that represents its entitlement.

We introduce the notion of a *plan repair market*, a plan repair mechanism based on the market paradigm. In each round r the problem owner o_r opens a auction in which actors may submit ask prices for the tasks of the candidates they are part of. Auctions in which sellers compete for the right to sell are called *procurement auctions*. We assume that agent have *private values*, i.e., an agent will not change its utility of a task once its learns the price submissions of other agents. This is a realistic assumption, as different agents will usually be submitting prices for different tasks.

We will use $q_{r,c} \in \mathbb{R}$ to denote the ask price of agent $a_{r,c}$ in round r for executing its part in candidate c. When $q_{r,c} = -u_{r,c}$ we say that $q_{r,c}$ is a *cost price*. For technical convenience, we also use notation

$$q_{i,j,r,c} = \begin{cases} q_{r,c} & \text{if } \ i = a_{r,c} \wedge j = o_r \\ 0 & \text{otherwise} \end{cases}$$

$q_{i,j,r,c}$ may be read as "the price that agent i asks to agent j for executing its task in candidate c in round r". We require $\forall i, j, r, c\colon\ u_{i,j,r,c} = 0 \rightarrow q_{i,j,r,c} = 0$. So, an agent may not ask a price other than 0 for a candidate in which it is not assigned a task. We use the following abbreviations.

$$q_{I,J,r,c} = \sum_{i \in I} \sum_{j \in J} q_{i,j,r,c}$$

$$q_{r,c} = q_{A,A,r,c}$$

$$q^{nett}_{i,r,c} = q_{A,i,r,c} - q_{i,A,r,c}$$

$q_{A,A,r,c}$ may be read as "the price asked by the actor in candidate c in round r to the owner in round r", i.e., $q_{a_{r,e_r},o_r,r,c}$. Furthermore, $q^{nett}_{i,r,c}$ may be read as "the price that will be paid by agent i in round r is c is elected". If $q^{nett}_{i,r,c}$ is negative, it denotes the price that agent i will *receive* in round r if candidate c is elected. Note that, for all agents i and rounds r, if c is a default candidate, $q^{nett}_{i,r,c} = 0$. We also use the following notation for the total income of an agent and the total expenditure of an agent respectively.

$$Q^{out}_{i,r} = \sum_{t=1}^{r} q_{i,A\setminus i,t,e_t} \qquad Q^{in}_{i,r} = \sum_{t=1}^{r} q_{A\setminus i,i,t,e_t}$$

Each agent has a *balance*, indicating its financial wealth. Let $b_{i,r} \in \mathbb{R}$ denote the balance of agent i in round r. In each round the balances are updated according to the payments that are made.

$$\forall i, r\colon\ b_{i,r} = \begin{cases} 0 & \text{if } \ r = 0 \\ b_{i,r-1} - q^{nett}_{i,r,e_r} & \text{if } \ r > 0 \end{cases} \tag{3}$$

This is under the assumption that a problem owner pays exactly the price asked by the actor in the elected candidate.

Lemma 3.1. $\forall i, r\colon\ b_{i,r} = Q^{out}_{i,r} - Q^{in}_{i,r}$

Proof. For $r = 0$, the proposition obviously holds. Suppose $b_{i,r} = Q^{out}_{i,r} - Q^{in}_{i,r}$ (IH). Then

$$\begin{aligned}
b_{i,r+1} &= b_{i,r} - q^{nett}_{i,r+1,e_{r+1}} \\
&= b_{i,r} + q_{i,A,r+1,e_{r+1}} - q_{A,i,r+1,e_{r+1}} \\
&= Q^{out}_{i,r} - Q^{in}_{i,r} + q_{i,A,r+1,e_{r+1}} - q_{A,i,r+1,e_{r+1}} && \text{(by (IH))} \\
&= Q^{out}_{i,r} - Q^{in}_{i,r} + q_{i,A\backslash i,r+1,e_{r+1}} + q_{i,i,r+1,e_{r+1}} - q_{A\backslash i,i,r+1,e_{r+1}} - q_{i,i,r+1,e_{r+1}} \\
&= Q^{out}_{i,r} + q_{i,A\backslash i,r+1,e_{r+1}} - Q^{in}_{i,r} - q_{A\backslash i,i,r+1,e_{r+1}} \\
&= Q^{out}_{i,r+1} - Q^{in}_{i,r+1}
\end{aligned}$$
□

When there is a minimum value below which balances are not allowed to go, we say that balances are *bounded.* We use $\beta \in \mathbb{R}^-$ to denote such a minimum value. We will shortly say more about this.

We assume for now that problem owners will want their problems to be solved against minimum costs and will therefore elect the cheapest candidate in any round. Formally, we assume (E1), defined as follows.

$$\text{(E1)} \quad : \quad \forall r\colon\ e_r = \min_c q_{r,c}$$

In this section we assume that agents ask cost prices. Formally, we assume (Q1), defined as follows.

$$\text{(Q1)} \quad : \quad \forall i, j, r, c\colon\ q_{i,j,r,c} = -u_{i,j,r,c}$$

Note that from this definition, it follows that a problem owner asks a price to itself for a default candidate, equal to the disutility of that candidate. This is merely a technical convenience, as electing the cheapest candidate will then result in electing the efficient candidate, given that agents ask cost prices and do not run into the bound. For this reason, $\forall r\colon\ q_{r,c} = q^{nett}_{o_r,r,c} = -u_{r,c}$ if c is not a default candidate and $\forall r\colon\ q_{r,c} = -u_{r,c} \ \wedge\ q^{nett}_{o_r,r,c} = 0$ if c is a default candidate. Furthermore, we assume that actors always perform the repair actions required from them by the elected candidate properly, such that the conflict of the problem owner is resolved.

We use in this and following sections a slightly different notation than in section 2. We used to write $U_{i,r}(\mathbf{e})$ to denote the cumulative utility of agent i in round r resulting from allocation $\mathbf{e}$. From now on, we will first assume an election rule, such as (E1), and a pricing rule, such as (Q1), and any other necessary preconditions, and then write $U_{i,r}$ to denote the cumulative utility of agent i in round r as a result of the allocation that occurs under the given preconditions. The same holds for Q^{out}, Q^{in}, b, and other variables we will introduce.

The following proposition shows that, in the plan repair market, if all agents always ask cost prices and have no budget constraints, the resulting allocation is efficient in every round.

Proposition 3.2 (Auctions)**.** *Under* (Q1)*,* (E1)*, if* $\beta = -\infty$*, for any* π*, the allocation* $\langle e_1, e_2, e_3, \ldots \rangle$ *is efficient.*

Proof. From (Q1) follows that $\forall r, c:\ q_{r,c} = -u_{r,c}$. From this and (E1), it follows that $\forall r:\ e_r = \max_c u_{r,c}$, thus, e_r is efficient for every round r. From proposition 2.1, it follows that the allocation $\langle e_1, e_2, e_3, \ldots \rangle$ is efficient. □

The following proposition shows that, if all agents ask cost prices and have no budget constraints, the balance of an agent is equal to the negation of its nett transferred utility.

Proposition 3.3. *Under* (Q1), (E1), *if* $\beta = -\infty$, *for any* π, *it holds that* $\forall i, r\colon\ b_{i,r} = -U^{nett}_{i,r}$.

Proof. For $r = 0$, the proposition obviously holds. Suppose $b_{i,r} = -U^{nett}_{i,r}$ (IH). It is easy to derive

$$\forall I, J, r, c\colon\ q_{I,J,r,c} = -u_{I,J,r,c} \qquad (*)$$

Then

$$\begin{aligned}
b_{i,r+1} &= b_{i,r} + q^{nett}_{i,r+1,e_{r+1}} \\
&= b_{i,r} + q_{i,A,r+1,e_{r+1}} - q_{A,i,r+1,e_{r+1}} \\
&= -U^{nett}_{i,r} - u_{i,A,r+1,e_{r+1}} + u_{A,i,r+1,e_{r+1}} && \text{(by (IH) and } (*)) \\
&= -(U^{out}_{i,r} - U^{in}_{i,r}) - (u_{i,A\backslash i,r+1,e_{r+1}} + u_{i,i,r+1,e_{r+1}}) + (u_{A\backslash i,i,r+1,e_{r+1}} + u_{i,i,r+1,e_{r+1}}) \\
&= -U^{out}_{i,r} - u_{i,A\backslash i,r+1,e_{r+1}} + U^{in}_{i,r} + u_{A\backslash i,i,r+1,e_{r+1}} \\
&= -U^{out}_{i,r+1} + U^{in}_{i,r+1} \\
&= -U^{nett}_{i,r+1}
\end{aligned}$$

□

To establish the relation between money and equity, we need to incorporate *scarcity* of money. The fact that humans do not have an infinite supply of money follows from the fact that banks impose lower bounds on the balances of its customers. In our model, $\beta \leq 0$ denotes the lower bound below which the balance of an agent is not allowed to go. Thus, it is forbidden to elect a candidate such that, after the payments have been made, $\exists i\colon\ b_{i,r} < \beta$.

The fact that balances are bounded can be expressed by defining a new election rule. We still assume that agents will want to elect the cheapest candidate possible, but only from the candidates they can afford. Formally, we assume (E2), defined as follows.

$$\text{(E2)} \quad : \quad \forall r\colon\ e_r = \min_c \{\, q_{r,c} \mid b_{o_r,r-1} - q^{nett}_{o_r,r,c} \geq \beta \,\}$$

The fact that balances are bounded provides an incentive for agents to earn money. If balances were unbounded, an agent would never have the need to earn money. Now, it must earn money in order for its balance to stay above the bound. If its balance is close to or equal to the bound, it can not afford to buy utility any more and will have to do without help, which is very inefficient.

Besides providing an incentive to collaborate, bounded balances also have an important effect on the equity of exchange. The following proposition shows that, if agents' balances are bounded and all agents ask cost prices, inequity of the resulting exchange is bounded by a constant value.

Proposition 3.4 (Bounded inequity)**.** *Let* π *denote* $\langle\, A, u, \mathbf{o}\,\rangle$ *and let* $k = |A|$. *Then, under* (Q1), (E2), *for all* π, β, *it holds that* $\forall r\colon\ U^-_r \geq (k-1)\cdot\beta$.

Proof. The proposition can be rewritten to $\forall r:\ \min\{U^{nett}_{1,r}, U^{nett}_{2,r}, \ldots, U^{nett}_{k,r}\} \geq (k-1)\cdot\beta$ with use of the definition of U^- (definition 3). Using proposition 3.3, the proposition can be rewritten to $\forall r\colon\ \max\{\, b_{1,r}, b_{2,r}, \ldots, b_{k,r}\,\} \leq (k-1)\cdot -\beta$. This is true since the maximal balance that an agent i can have is achieved when all agents have spent all their credits on i, amounting to $(k-1)\cdot -\beta$. □

Although bounded balances ensure a bound on inequity, the resulting allocations need not be efficient anymore.

Proposition 3.5 (Inefficiency)**.** *Under* (Q1)*,* (E2)*, if* $\beta \neq -\infty$ *then the exchange may be inefficient.*

Proof. Suppose that in some round r, agent 1 is problem owner and has balance $b_{1,r} = \beta$. There are two repair candidates, $c_1 = \langle -100, 0 \rangle$ and $c_2 = \langle 0, -10 \rangle$. Agent 1 now can not afford c_2 and has to resort to the inefficient c_1. □

Thus, it might happen that in a certain round the problem owner has not enough money to buy the cheapest candidate. This can be justified from the viewpoint of equity, since this problem owner has already received a lot of utility from others and should at this point provide utility before receiving more. It must elect default candidates then, i.e., help itself, until it has sufficient funds again. Also, it could offer tasks below cost price to earn extra cash.

The market mechanism with bounded balances will achieve a certain trade-off between efficiency and equity, depending on the set of occurring problems and repair candidates and the height of the bound. For instance, if $\beta = 0$, only default candidates can be elected. This would be highly inefficient but perfectly equitable. If however $\beta = -\infty$, efficient candidates will be elected in every round but inequity is unbounded. By setting the bound somewhere in between, a mechanism designer, or the participants themselves, can determine the trade-off that is appropriate for their particular exchange scenario.

Note that inefficiency is often the result of introducing fairness constraints. In many scenario's, allocating resources both efficient and fair is not possible [5, 27, 20] and a trade-off between the two has to be found (cf. [12]).

4 Selfish agents

Until now we have assumed that agents submit cost prices for their tasks. In this section we investigate what happens if agents ask prices possibly above their cost price. A very normal phenomenon in markets is that sellers will try to set their ask prices just below that of the competition. This leads to a high demand for the product on sale and, if the selling price is above the cost price, results in a profit on every item sold.

From the viewpoint of social exchange, this behaviour is a little less ordinary. We argued in the previous section that money is a distributed system for reciprocal exchange. If agents over-price their goods or services, the direct correspondence between prices and utility is lost. An agent that over-prices its good or service may obtain a sum of money that it can spend to obtain an amount of utility possibly larger than what it provided. In this way, the resulting exchange may not have the desired equity properties described in the previous section.

Especially when an agent is dependent on another agent, the difference between price and utility may be large. Note that in many real-life market situations, dependency between buyer and seller is weak, as there are a lot of other sellers (and buyers) for a given product. The competition between these sellers drives prices down and sellers will try to minimize their production costs to maximize their profit margin. As a result, prices will often lie close together and not far above production costs.

In the plan repair scenario however, competition is sometimes not so strong, leading to situations with a strong dependence of the buyer on the seller. Consider the following scenario,

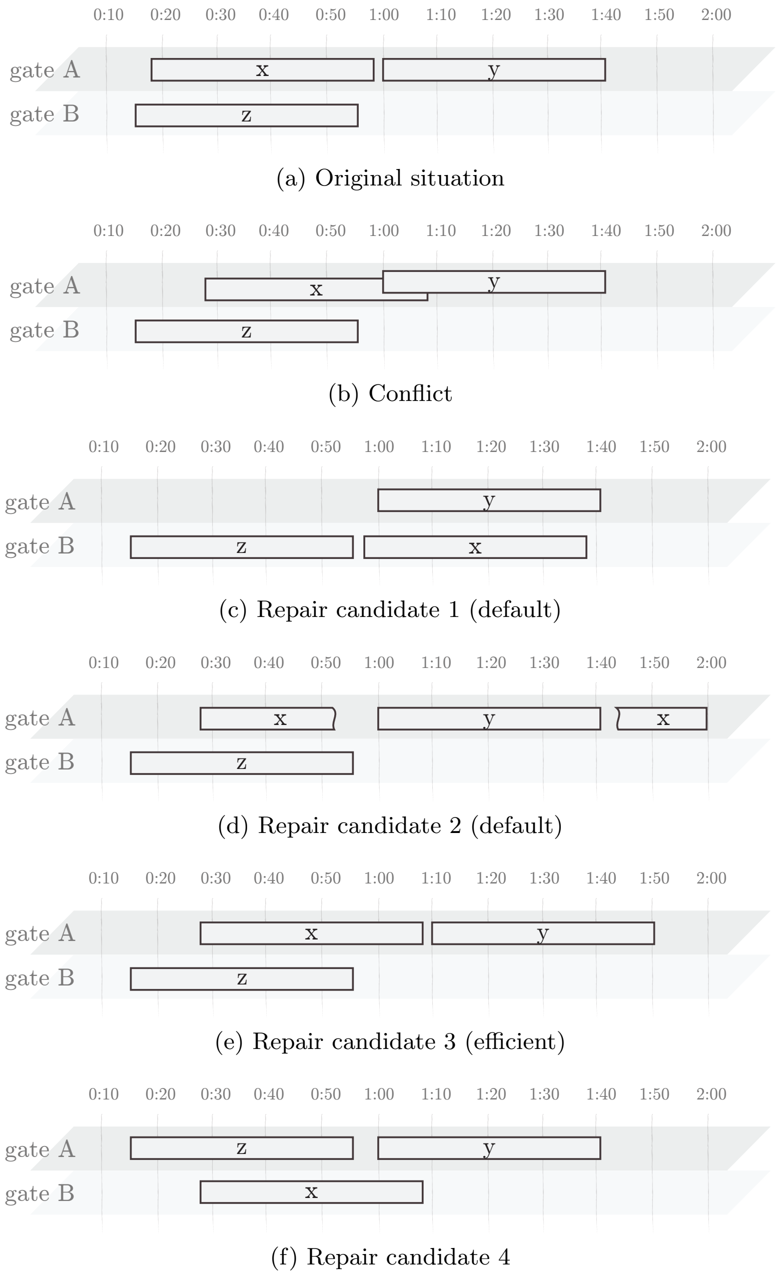

(a) Original situation

(b) Conflict

(c) Repair candidate 1 (default)

(d) Repair candidate 2 (default)

(e) Repair candidate 3 (efficient)

(f) Repair candidate 4

Figure 1: Planning conflict and repair candidates.

depicted in figure 1. An arriving aircraft x is delayed and will enter its gate ten minutes later. As a result it wants to leave it ten minutes later than planned (figure 1b). For this it needs the cooperation of aircraft y, scheduled next, which should wait ten minutes (figure 1e). This would require very little effort of y and there is no one scheduled after y. Other repair candidates for x involve gate changes (figures 1c and 1f) or going off-gate prematurely (figure 1d), and all have a much lower utility for x or other aircraft. Thus, the repair candidate involving y is the obvious choice. Aircraft y now realizes its strategic advantage and can ask a price far above its disutility, just below the second lowest price. Its candidate is still the cheapest and will thus be chosen, providing aircraft y with an attractive profit and aircraft x with a very expensive deal.

We will call the phenomenon of asking a price just below that of the second cheapest seller *exploitation*. From a social viewpoint, the cheapest seller exploits the dependency that exist between him and the buyer. If this dependency is strong, i.e., if the difference in price between the two cheapest candidates is large, the resulting profit as well as the disadvantage incurred by the buyer can be large.

Exploitation can have a significant effect on the equity of exchange. If an agent is structurally able to exploit, the profits thus derived can be used to obtain utility from others. In this way an exploiter may be able to structurally obtain more utility than it provides to the community, resulting in an unfair advantage, and to inequitable exchange.

We will now formally define exploitation in the plan repair market and show its dominance over not exploiting and the possible effect it can have on equity of exchange. We assume that some agents, *exploiters*, will exploit if the possibility exists, and other agents, the *coalition agents*, will, in principle, ask cost prices. Let σ^c be the coalition agent type and σ^e the exploiter agent type. An *agent type vector* $\boldsymbol{\sigma} = \langle \sigma_1, \sigma_2, ..., \sigma_k \rangle$ specifies for each agent i its type $\sigma_i \in \{ \sigma^c, \sigma^e \}$. Let $\tilde{\mathbf{c}}_\mathbf{r} = \langle \tilde{c}_r^1, \tilde{c}_r^2, \ldots, \tilde{c}_r^m \rangle$ be the vector of candidates in round r ordered descendingly by utility, i.e., the vector $\langle 1, 2, \ldots, m \rangle$ in which the elements are ordered such that $u_{r,\tilde{c}_r^1} \geq u_{r,\tilde{c}_r^2} \geq \ldots \geq u_{r,\tilde{c}_r^m}$. We assume that an exploiter will, if it is actor in the cheapest candidate, ask a price just below that of the second cheapest seller. Formally, we assume

$$(\text{Q2}) \quad : \quad \forall i, j, r, c\colon \; q_{i,j,r,c} = -u_{i,j,r,c} + q^e_{i,j,r,c}$$

where

$$q^e_{i,j,r,c} = \begin{cases} -(u_{r,\tilde{c}_r^2} - u_{r,\tilde{c}_r^1}) - \epsilon & \text{if } \; \sigma_i = \sigma^e \;\wedge\; i = a_{r,c} \;\wedge\; j = o_r \;\wedge\; i \neq j \;\wedge\; c = \tilde{c}_r^1 \\ & \qquad \wedge\; u_{r,\tilde{c}_r^2} < u_{r,\tilde{c}_r^1} \\ 0 & \text{otherwise} \end{cases}$$

Thus, an exploiter will raise its price until just below that of the second most efficient candidate, if it is actor in the efficient candidate, it is not problem owner and the second most efficient candidate is strictly less efficient.

The following proposition shows that the presence of exploiters may lead to an ever increasing inequity of exchange.

Proposition 4.1 (Inequity under exploitation)**.** *Let* π *denote* $\langle A, u, \mathbf{o} \rangle$ *and let* $k = |A|$*. Let* $\boldsymbol{\sigma}$ *be a type vector such* $\sigma_i = \sigma^e$ *for some agent* i*. Then, under* (Q2), (E2)*, there exist* π *and* β *such that* $\forall r\colon U_r^- \leq (r-2) \cdot \beta$*.*

Proof. We show that inequity of exchange after r rounds can be as high as $(r-1)\cdot\beta$. We do this by constructing an example in which an exploiter alternates between exploiting maximally and spending maximally. The example is depicted in figure 2.

Let $A = \{ x, y \}$ and $\sigma_y = \sigma^e$. Suppose that $o_1 = x$, $a_{1,e_r} = y$ and $u_{1,e_r} = -\epsilon$. Thus, agent

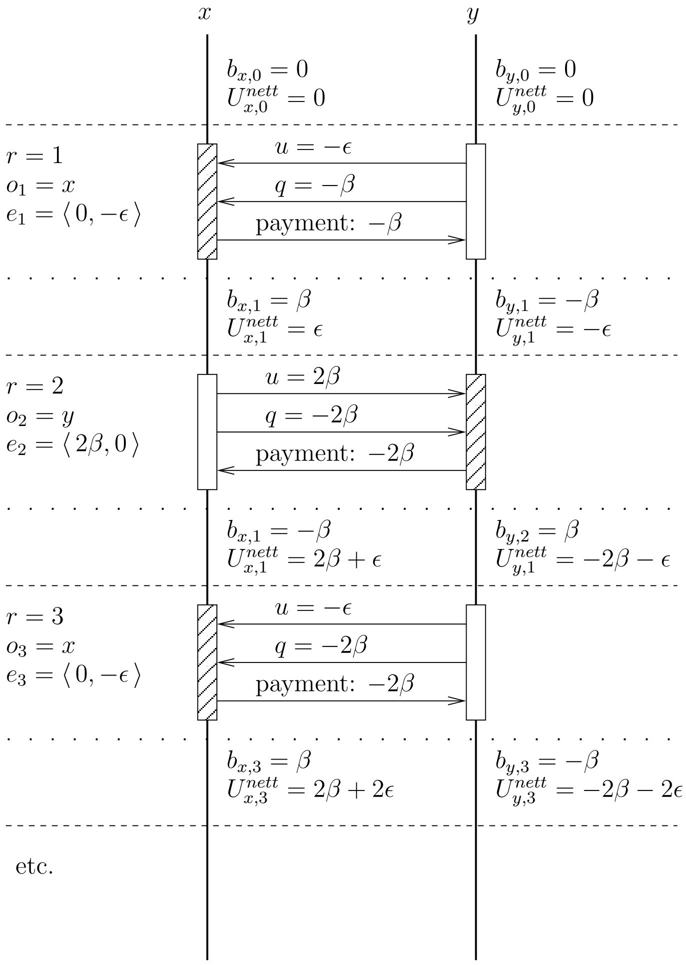

Figure 2: Inequity as a result of exploitation. Shaded bar denotes problem owner.

y is almost indifferent between performing the action or not. Agent y, being an exploiter, asked price $-u_{r,\tilde{c}^2_r} - \epsilon$. Suppose this price is equal to $-\beta$. Then, after enforcing the candidate and making the payment, $b_x = \beta$, $b_y = -\beta$, $U^{nett}_{x,1} = \epsilon$ and $U^{nett}_{y,1} = -\epsilon$.

Suppose that $o_2 = y$, $a_{2,e_r} = x$, $u_{2,e_r} = 2 \cdot \beta$ and $q_{2,e_r} = -2 \cdot \beta$. After enforcing this candidate en making the corresponding payments, $b_x = -\beta$, $b_y = \beta$, $U^{nett}_{x,2} = 2 \cdot \beta + \epsilon$, $U^{nett}_{y,2} = -2 \cdot \beta - \epsilon$.

Suppose that $o_3 = x$, $q_{3,e_r} = y$ and $u_{3,e_r} = -\epsilon$. Suppose $q_{3,e_r} = -2 \cdot \beta$. Agent x can pay the price since $b_x = -\beta$. After enforcing the candidate and making the payment, $b_x = \beta$, $b_y = -\beta$, $U^{nett}_{x,3} = 2 \cdot \beta + 2\epsilon$, $U^{nett}_{y,3} = -2 \cdot \beta - 2\epsilon$.

Suppose that rounds two and three are repeated infinitely. In every even round after round 3, agent x performs a task of utility $2 \cdot \beta$, for which it receives $-2 \cdot \beta$, the nett transferred utility of x changes with $2 \cdot \beta$ and that of y changes with $-2 \cdot \beta$. In every uneven round after round 3, agent y performs a task of utility $-\epsilon$, for which it receives $-2 \cdot \beta$, the nett transferred utility of agent x changes with ϵ and that of agent y with $-\epsilon$. This gives, for $r > 3$, $U^{nett}_{x,r} = \left\lfloor \frac{r}{2} \right\rfloor \cdot 2 \cdot \beta + \left\lfloor \frac{r-1}{2} \right\rfloor \cdot 2\epsilon$, which is less than $(r-2) \cdot \beta$. $\square$

To be able to compare strategies, we need a measure that indicates an agent's well-being at a certain moment. This measure should depend on two factors: the cumulative utility of an agent and its balance. The cumulative utility $U_{i,r}$ is important since it indicates the total amount of effort that agent i had to expend until round r. The height of an agent's balance $b_{i,r}$ is important for the following reason. In the scenario with bounded balances, every agent will have some ideal value they want their balance to be equal to, that is, if it is lower than that value they consider the risk of running into the bound too high and will want to clear their debts, and if it is higher than that value they feel less urge to help others, as they already have enough money and are 'safe'. This value could be different for different agents at different moments in time, but for simplicity we assume this value zero for all agents at each moment. Thus, an agent with a negative balance wants to clear this debt and expects a utility decrease as a result. An agent with a positive balance may spend this without consequences and expects a positive result on its utility (actually lack of a negative effect on its utility compared to when its balance would be lower).

We define the *perceived utility* $U^p_{i,r}$ of agent i in round r to be

$$U^p_{i,r} = U_{i,r} + b_{i,r} \tag{4}$$

Note that defining agents' utilities in this way is known as assuming that agents have *quasi-linear* utilities[2]. Also, we introduce the notion of a *scenario*, which is a tuple $\langle \pi, \boldsymbol{\sigma}, \beta \rangle$ consisting of a problem set π, a type vector $\boldsymbol{\sigma}$ and a bound β. Cumulative utilities, balances, perceived utilities, etc. are defined with respect to a scenario, a pricing rule and an election rule. We use $U_{i,r}(\zeta)$ to denote the cumulative utility of agent i in round r in scenario ζ under some given pricing rule and election rule. $U^p_{i,r}(\zeta)$, $b_{i,r}(\zeta)$ and other variables go in similar fashion.

The following proposition shows that exploitation is a dominant strategy over not exploiting. We prove it for the case where no agent runs into its bound, i.e., the bound never influences the winning candidate.

Theorem 4.2 (Dominance of exploitation strategy)**.** *Let $\boldsymbol{\sigma}$ be a type vector such that $\sigma_i = \sigma^c$ for some agent i, and let $\boldsymbol{\sigma'}$ be a type vector such that $\boldsymbol{\sigma'} = \langle \sigma'_1, \sigma'_2, \ldots, \sigma'_k \rangle$, $\sigma'_i = \sigma^e$ and $\forall j : j \neq i \rightarrow \sigma_j = \sigma'_j$. Let t be a round such that i has had the opportunity to exploit in some round $s < t$. Then, under* (Q2), (E2)*, for any π, we have that $U^p_{i,t}(\langle \pi, \boldsymbol{\sigma}, -\infty \rangle) < U^p_{i,t}(\langle \pi, \boldsymbol{\sigma'}, -\infty \rangle)$.*

---

[2] One could argue that, as an agent is in the end interested in its cumulative utility only, the perceived utility should be defined as $U^p_{i,r} = U_{i,r} + X_{i,r}$, where $X_{i,r}$ is the effect of having $b_{i,r}$ on i's future cumulative utility. This however leads to a paradoxical situation. Consider the following situation. Suppose that all agents are coalition agents and one agent x has $b_x = 0$. In the rounds to come, he will want to receive services with a total disutility of 100. As he wants his balance to be equal to 0, he will want to earn 100 credits, so he will have to perform services with disutility 100. Now suppose that he had $b_x = b$. In that case he would have needed $100 - b$ credits, so he would have had to perform services with disutility $100 - b$. So, having a balance of b instead of 0 reduces the disutility of services to perform with b. Thus, $X_{x,r} = b_{x,r}$.

Now suppose that x is an exploiter. It has $b_x = 0$ and needs 100 credits in the coming rounds. To earn this, he will have to perform services of total disutility less than 100, since it makes profits from exploitation. Say that he makes an average profit of 10%. He will have to perform services with total disutility of $100/1.1 = -90.9$. Now suppose that he had $b_x = b$. In that case he would have needed $100 - b$ credits, so he would have had to perform services with disutility $(100 - b)/1.1 = 90.9 - b/1.1$. So, having a balance of b instead of 0 reduces the disutility of services to perform with $b/1.1$. Thus, $X_{x,r} \leq b_{x,r}$.

Thus, between two agents with equal, positive balances and equal cumulative utilities, if one is exploiter and the other coalition agent, the exploiter will have a lower perceived utility, which is counter-intuitive.

Proof. Let $\zeta = \langle \pi, \boldsymbol{\sigma}, -\infty \rangle$ and $\zeta' = \langle \pi, \boldsymbol{\sigma'}, -\infty \rangle$. To shorten notation a bit, we write $U_{i,r}$, $Q^{in}_{i,r}$, e_r, etc. to denote $U_{i,r}(\zeta)$, $Q^{in}_{i,r}(\zeta)$, $e_r(\zeta)$, etc., and write $U'_{i,r}$, $Q^{in\,\prime}_{i,r}$, e'_r, etc. to denote $U_{i,r}(\zeta')$, $Q^{in}_{i,r}(\zeta')$, $e_r(\zeta')$, etc..

The fact that i exploits in ζ' does not influence the elected candidate in any round, i.e., $\forall r: \; e_r = e'_r$. As a result, $\forall r: \; U_{i,r} = U'_{i,r}$. Also, each agent other than i follows the same strategy in ζ and ζ'. Therefore, each agent other than i will ask the same prices in ζ and ζ'. As a result, $\forall r\colon\, Q^{in}_{i,r} = Q^{in\,\prime}_{i,r}$.

From (Q2) follows that $\forall r: \; q_{i,A,r,e_r} \leq q'_{i,A,r,e_r}$. In round s, we have $q_{i,A,s,e_s} < q'_{i,A,s,e_s}$. From these two facts we may conclude $\forall r \geq s\colon\, Q^{out}_{i,r} < Q^{out\,\prime}_{i,r}$. For $r \geq s$, in particular round t, we can now rewrite

$$\begin{aligned} U^p_{i,r} &= U_{i,r} + b_{i,r} \\ &= U_{i,r} + Q^{out}_{i,r} - Q^{in}_{i,r} && \text{(by lemma 3.1)} \\ &< U'_{i,r} + Q^{out\,\prime}_{i,r} - Q^{in\,\prime}_{i,r} \\ &= U^{p\,\prime}_{i,r} \end{aligned}$$

□

For unbounded balances, dominance of exploitation is easily proven. With bounded balances, exploitation is not necessarily dominant as an example problem set can be constructed in which a coalition agent is better off than an exploiter. However, such problem sets are highly construed. We hypothesise that in most instances, exploiters will perform better than coalition agents under bounded balances. We prove this hypothesis by experiment in section 7.

As exploitation is a better performing strategy, agents can be expected to adopt this strategy. If some agents exploit, inequity is no longer bounded. In section 7 we will, by means of experiment, assess the level of inequity that will occur in a realistic ATM scenario.

4.1 Large versus small airlines

Exploitation does not necessarily lead to inequity. If, in a hypothetical situation, all agents would gain as much from exploiting as they would loose from being exploited, equity is maintained. In the airport plan repair scenario however, full scale exploitation is likely to lead to significant inequity. This follows from the fact that airlines operating on an airport usually greatly differ in size. Most major airports have a hub-owner who operates the majority of flights. Besides the hub-owner there are a number of other airlines, from medium-sized players to very small ones.

A large airline operates many flights, and will therefore more often run into planning conflicts. It will thus need the collaboration of other airlines often. At the same time, its collaboration will be asked for by other often too. Thus, its nett transferred utility and its balance will remain stable over time if agents are honest. Being large does not give an immediate advantage in the cooperative scenario.

When agents may exploit however, being a large airline has an important consequence: a large airline will typically have better default candidates. As a large airline has many flights in many slots on many gates and runways, it will have more options to solve the conflict by itself. It may for instance delay one of its flights, swap two slots or rebook passengers to another of its flights. Therefore, its best default candidate will on average have a higher utility than that of a small airline. This has an effect on the amount to which it can be exploited. Having an attractive default candidate means that exploiters in other candidates can raise their price less (or not at all) than they would have if the default candidate was worse. Thus, a large airline is less 'exploitable' than a small airline.

Also, a large airline is a better exploiter than a small airline. This is the result of a phenomenon we have not described yet. If an exploiter is actor in the cheapest candidate, but it is actor in the second-to-cheapest candidate too, it can raise its price in both candidates. In fact, it can raise its price in all the candidates in which it is actor until just below the price of the cheapest candidate in which it is not actor. In that way the exploiter boosts its income. A large airline will be actor in candidates more often than a small airline. Therefore, the possibility that it is actor in two or more cheapest candidates will be larger. As a result, it will relatively gain more profits from exploitation than a small airline.

Raising ones price in more than one candidate is a strategy which we have not involved in the theory in order to keeps its complexity to a minimum. We have implemented it however in the experiments in section 7.

Thus, the negative consequences of exploitation are likely to occur in the airport plan repair scenario. Large airlines are expected to gain a significant advantage over small airlines. For that reason, applying the market mechanism with standard currency to the airport plan repair problem will lead to undesired inequity.

4.2 Swindling

Besides raising the ask price, an agent can also lessen its disutility when providing a service to make a profit. We can call this *swindle*, as the seller does not provide the promised or required good or service (cf. [17]). Swindle is equivalent to exploitation in the sense that the amount of money paid for the good or service is too high compared to the quality of the provided good or service. The difference is that exploitation reveals itself early, a buyer willingly pays the ask price, while swindle reveals itself late, a buyer inspecting its good or receiving the service notices that it is below expectation. Also, exploitation is legally allowed in most market scenario's, while swindle is often not.

In the airport plan repair scenario swindle would happen if an airline does not fully perform the actions specified in the repair candidate. The results of this can be substantial as plannings are often tight. For instance, a flight might miss its departure slot as a result of a few minutes delay, or transfer passengers might not be able to board their second flight, leading to forced rebookings.

To prevent oneself against swindle, participants normally enter into a contract specifying the good or service that is to be provided and measures that will be taken if the good or services does not meet the given conditions.

If one is the victim of swindle nevertheless, for instance because of an unforseen shortcoming or because a contract was not made, the result will usually be a decreased credibility of the seller in the eyes of the buyer. In the plan repair scenario, this may lead to the buyer not wanting to do business with the seller anymore, out of fear of being deceived again.

Exploitation and swindling are both forms of what we call *under-delivering*; providing less than promised. In fact, the measures we take against exploitation in future sections may be used as measures against under-delivering in general. But, as exploitation is much less easily counteracted than swindling, and because our focus on exploitation is relatively original, we will continue to use that term.

5 Enforcing equity

In this section we investigate whether agents can counteract exploitation by means of adapting ask prices. More precisely, we investigate whether agents can use a pricing scheme such that exploiters are penalized and exploiting becomes unattractive. Thus, we aim for an incentive-compatible plan repair mechanism, in which the collaborative strategy is dominant. If this can be achieved, agents will want to adopt the collaborative strategy of asking cost prices, which, as we saw in previous sections, leads to an exchange with desirable levels of efficiency and equity.

We make the – rather optimistic – assumption that a problem owner that is being exploited is exactly able to observe the level of exploitation. Thus, it knows both the utility of the task performed and the price asked. Suppose that all agents communicate this information to each other and thus know at any moment the amount of money that a certain exploiter y has earned by exploiting. Suppose that all agents add this value as a penalty to the prices they ask to exploiter y. In that way an exploiter would be deprived of its exploitation profits.

Note that this penalty system is the collaborative equivalent of the well-known tit-for-tat strategy. An agent that has over-prices its service will find that all other agents over-price their services in return.

Let $q^p_{i,j,r,c}$ denote the *penalty* that agent i charges to agent j in round r for candidate c. Also, let

$$q^e_{I,J,r,c} = \sum_{i\in I}\sum_{j\in J} q^e_{i,j,r,c} \qquad\qquad q^p_{I,J,r,c} = \sum_{i\in I}\sum_{j\in J} q^p_{i,j,r,c}$$

and

$$Q^e_{I,J,r} = \sum_{t=0}^{r} q^e_{I,J,t,e_t} \qquad\qquad Q^p_{I,J,r} = \sum_{t=0}^{r} q^p_{I,J,t,e_t}$$

We assume from now on that agents penalize exploiters. Formally, we assume (Q3), defined as follows.

$$\text{(Q3)} \quad : \quad \forall i,j,r,c\colon\; q_{i,j,r,c} = -u_{i,j,r,c} + q^p_{i,j,r,c} + q^e_{i,j,r,c}$$

where

$$q^p_{i,j,r,c} = \begin{cases} Q^e_{j,A,r-1} - Q^p_{A,j,r-1} + \epsilon & \text{if } \; i = a_{r,c} \;\wedge\; j = o_r \;\wedge\; i \neq j \;\wedge\; Q^e_{j,A,r-1} - Q^p_{A,j,r-1} > 0 \\ 0 & \text{otherwise} \end{cases}$$

and $q^e_{i,j,r,c}$ is as defined before. Thus, agents add to the prices they ask to an exploiter (other than themselves) an amount equal to the income made by this exploiter by exploiting minus the penalties that he already had to pay for it, plus a small amount ϵ. The extra small amount ϵ is to make sure that exploiters, after exploiting and being retaliated, are worse off than if they had not exploited. Note that $\forall i,j,c:\; q^e_{i,j,0,c} = q^p_{i,j,0,c} = 0$ and thus $\forall j:\; Q^e_{j,A,0} = Q^p_{A,j,0} = 0$. Also note that exploiters charge penalties too. We assume that exploiters, although guilty of exploitation themselves, will be willing to impose penalties on other exploiters as they themselves are victim of other exploiters too and it is in fact advantageous to charge penalties.

Until now, we assumed that electing the cheapest candidate was the rational thing to do for a problem owner. With penalty pricing in effect, this is no longer the case. When an exploiter is problem owner, the other agents might raise their price and this might have the result that a default candidate becomes the cheapest candidate. In that case, the exploiter should not elect this default candidate. If it would, other agents would simply add the penalty again the next time the exploiter is problem owner. The other agents will continue to charge the penalty until the exploiter has paid it. As a result, it does not matter whether he pays it directly or later. Thus, the exploiter should not let his choice of candidate depend on the penalty.

We assume (E3) from now on, defined as follows.

$$\text{(E3)} \quad : \quad \forall r\colon\ e_r = \min_c \{\, q_{r,c} - q^p_{r,c} \mid b_{o_r,r-1} - q^{nett}_{o_r,r,c} \geq \beta \,\}$$

Thus, a problem owner selects the cheapest candidate *as if penalties were not charged*[3].

Imposing penalties takes away the advantage of an exploiter. This can be seen most clearly in a scenario with one exploiter; its nett transferred utility can not grow infinitely anymore. To prove this, we first need to prove the following equality.

Proposition 5.1. $\forall i, r\colon\ b_{i,r} = -U^{nett}_{i,r} + Q^e_{i,A,r} - Q^e_{A,i,r} + Q^p_{i,A,r} - Q^p_{A,i,r}$

Proof. For round 0, the proposition obviously holds. Suppose $b_{i,r-1} = -U^{nett}_{i,r-1} + Q^e_{i,A,r-1} - Q^e_{A,i,r-1} + Q^p_{i,A,r-1} - Q^p_{A,i,r-1}$ (IH). We may derive

$$\begin{aligned}
b_{i,r} &= b_{i,r-1} + q_{i,A,r,e_r} - q_{A,i,r,e_r} && \text{(by (3))}\\
&= -U^{nett}_{i,r-1} + Q^e_{i,A,r-1} - Q^e_{A,i,r-1} + Q^p_{i,A,r-1} - Q^p_{A,i,r-1} + \\
&\quad - u_{i,A,r,e_r} + q^p_{i,A,r,e_r} + q^e_{i,A,r,e_r} - (-u_{A,i,r,e_r} + q^p_{A,i,r,e_r} + q^e_{A,i,r,e_r}) && \text{(by (IH) and (Q3))}\\
&= (-U^{nett}_{i,r-1} - u_{i,A,r,e_r} - -u_{A,i,r,e_r}) + (Q^e_{i,A,r-1} + q^e_{i,A,r,e_r}) - \\
&\quad (Q^e_{A,i,r-1} + q^e_{A,i,r,e_r}) + (Q^p_{i,A,r-1} + q^p_{i,A,r,e_r}) - (Q^p_{A,i,r-1} + q^p_{A,i,r,e_r})\\
&= -U^{nett}_{i,r} + Q^e_{i,A,r} - Q^e_{A,i,r} + Q^p_{i,A,r} - Q^p_{A,i,r}
\end{aligned}$$

□

The following proposition shows that under penalty pricing, if there is only one exploiter, its nett transferred utility is bounded by a constant value.

Proposition 5.2 (Bounded inequity)**.** *Let* $\boldsymbol{\sigma}$ *be a type vector such that* $\sigma_i = \sigma^e$ *for some agent* i *and that* $\forall j:\ j \neq i \rightarrow \sigma_j = \sigma^c$*. Then, under* (Q3), (E3), *for all* π, β, *it holds that* $U^{nett}_{i,r} \leq -\beta$.

Proof. As i is the only exploiter, we have $\forall r\colon\ Q^e_{A,i,r} = Q^p_{i,A,r} = 0$. Filling this in in proposition 5.1 gives $\forall r\colon\ b_{i,r} = -U^{nett}_{i,r} + Q^e_{i,A,r} - Q^p_{A,i,r}$ $(*)$. Then, observe that $U^{nett}_{i,r}$ can only increase in rounds where i is problem owner. Let t be a round in which $o_t = i$. Then, $q^e_{i,A,t,e_t} = 0$. We have

$$\begin{aligned}
Q^e_{i,A,t} - Q^p_{A,i,t} &= (Q^e_{i,A,t-1} + q^e_{i,A,t,e_t}) - (Q^p_{A,i,t-1} + q^p_{A,i,t,e_t})\\
&= Q^e_{i,A,t-1} - (Q^p_{A,i,t-1} + q^p_{A,i,t,e_t})
\end{aligned}$$

[3] Alternatively, we could have used (E2) with (Q3) defined such that the problem owner also added $Q^e_{j,A,r-1} - Q^p_{A,j,r-1} + \epsilon$ to the costs of its default candidates.

Suppose that $Q^e_{i,A,t-1} - Q^p_{A,i,t-1} \leq 0$. According to (Q3), $q^p_{A,i,t,e_t} = 0$ and it follows that $Q^e_{i,A,t} - Q^p_{A,i,t} \leq 0$.

Now suppose that $Q^e_{i,A,t-1} - Q^p_{A,i,t-1} > 0$. According to (Q3), $q^p_{A,i,t,e_t} = Q^e_{j,A,t-1} - Q^p_{A,j,t-1} + \epsilon$. It follows that $Q^e_{i,A,t} - Q^p_{A,i,t} = -\epsilon$.

Thus, $Q^e_{i,A,t} - Q^p_{A,i,t} \leq 0$. Filling this in in $(*)$ gives $b_{i,t} \leq -U^{nett}_{i,r}$. As $b_{i,t} \geq \beta$, we derive $U^{nett}_{i,t} \leq -\beta$ □

The following theorem shows that under penalty pricing, an exploiter is equal or worse off than a coalition agent in each round in which it is problem owner, if balances are unbounded.

Theorem 5.3 (Dominance of coalition strategy). *Let $\boldsymbol{\sigma}$ be a type vector such that $\sigma_i = \sigma^c$ for some agent i, and let $\boldsymbol{\sigma'}$ be a type vector such that $\sigma'_i = \sigma^e$ and $\forall j\colon j \neq i \rightarrow \sigma_j = \sigma'_j$. Then, under* (Q3), (E3)*, for all π, in each round r in which $o_r = i$, it holds that $U^p_{i,r}(\langle \pi, \boldsymbol{\sigma}, -\infty \rangle) \geq U^p_{i,r}(\langle \pi, \boldsymbol{\sigma'}, -\infty \rangle)$.*

Proof. Let $\zeta = \langle \pi, \boldsymbol{\sigma}, -\infty \rangle$ and $\zeta' = \langle \pi, \boldsymbol{\sigma'}, -\infty \rangle$. We write $U_{i,r}$, $Q^{in}_{i,r}$, e_r, etc. to denote $U_{i,r}(\zeta)$, $Q^{in}_{i,r}(\zeta)$, $e_r(\zeta)$, etc., and write $U'_{i,r}$, $Q^{in\,\prime}_{i,r}$, e'_r, etc. to denote $U_{i,r}(\zeta')$, $Q^{in}_{i,r}(\zeta')$, $e_r(\zeta')$, etc..

From (E3), it follows that the elected candidate in a round does not depend on the penalty, i.e., it depends on $-u_{r,c} + q^e_{r,c}$. As we know that exploitation never influences the elected candidate, we may conclude that the elected candidate in a round only depends on $-u_{r,c}$, i.e., on $\mathfrak{u}$. As ζ and ζ' contain the same problem set π and thus the same utility profile $\mathfrak{u}$, we may conclude that $\forall r\colon e_r = e'_r$. From this we may conclude $\forall r\colon U^{nett}_{i,r} = U^{nett\,\prime}_{i,r}$.

From the definition of (Q3), it follows that the amount to which on is exploited by others does not depend on one's type. The same holds for the amount of penalties one asks to others. Thus, we have $Q^e_{A,i,r} = Q^{e\,\prime}_{A,i,r}$ and $Q^p_{i,A,r} = Q^{p\,\prime}_{i,A,r}$. Since i does not exploit in ζ, we have $Q^e_{i,A,r} = 0$ and $Q^p_{A,i,r} = 0$. Using proposition 5.1, we obtain

$$\begin{aligned}
b_{i,r} &= -U^{nett}_{i,r} + Q^e_{i,A,r} - Q^e_{A,i,r} + Q^p_{x,A,r} - Q^p_{A,x,r} \\
&= -U^{nett}_{i,r} - Q^e_{A,i,r} + Q^p_{i,A,r} \\
&= -U^{nett\,\prime}_{i,r} - Q^{e\,\prime}_{A,i,r} + Q^{p\,\prime}_{i,A,r} \\
&= (-U^{nett\,\prime}_{i,r} + Q^{e\,\prime}_{i,A,r} - Q^{e\,\prime}_{A,i,r} + Q^{p\,\prime}_{i,A,r} - Q^{p\,\prime}_{A,i,r}) - Q^{e\,\prime}_{i,A,r} + Q^{p\,\prime}_{A,i,r} \\
&= b'_{i,r} - Q^{e\,\prime}_{i,A,r} + Q^{p\,\prime}_{A,i,r}
\end{aligned}$$

From $\forall r\colon e_r = e'_r$ we may also conclude $\forall r\colon U_{i,r} = U'_{i,r}$. In the proof of proposition 3.4, we showed that $Q^{e\,\prime}_{i,A,r} - Q^{p\,\prime}_{A,i,r} \leq 0$ for rounds in which $o'_r = i$. Thus we may derive

$$\begin{aligned}
U^p_{i,r} &= U_{i,r} + b_{i,r} \\
&= U'_{i,r} + b'_{i,r} - Q^{e\,\prime}_{i,A,r} + Q^{p\,\prime}_{A,i,r} \\
&\geq U^{p\,\prime}_{i,r}
\end{aligned}$$
□

In section 7, we prove by experiment the expected dominance of the coalition strategy under penalty pricing under bounded balances. The penalty pricing rule manages to make exploitation unattractive. As a result, agents will want to be coalition agents. If all agents are coalition agent, the desired trade-off between efficiency and equity results.

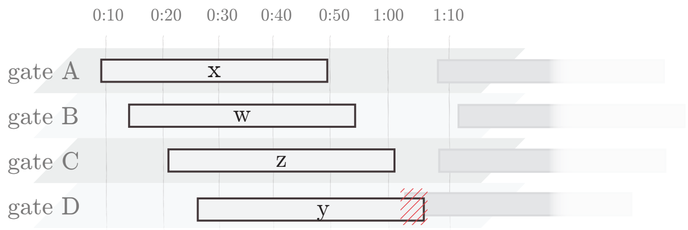

(a) Conflict

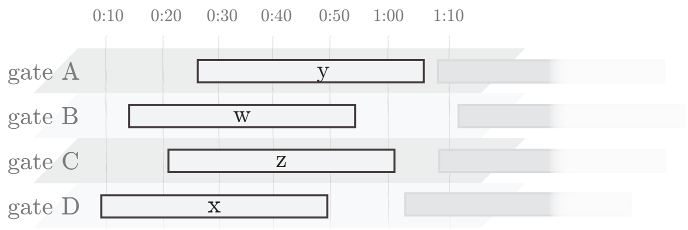

(b) Repair candidate 1

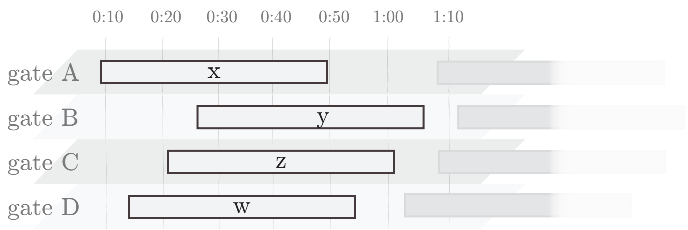

(c) Repair candidate 2

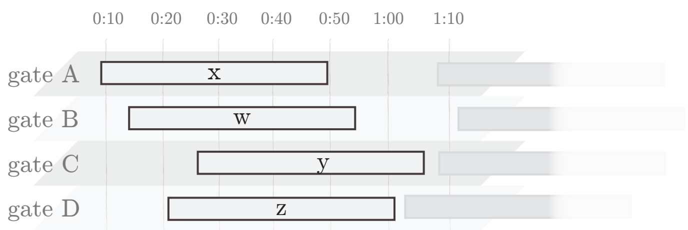

(d) Repair candidate 3

Figure 3: A scenario in which agents can forsake.

5.1 Forsaking

Although penalization as described in the previous section effectively takes away the advantage of exploitation, it is not necessarily the best strategy for agents to follow. In a situation where *competition* between coalition agents occurs, it can be very profitable for agents to deviate from the coalition pricing rule. Consider the situation depicted in figure 3. Agent y has run into a conflict with another agent, as can be seen in figure 3a. Somehow this other agent is not able to shift its reservation to make room for y. Three other candidates are possible, shown in figures 3b, 3c and 3d. Each of these candidates involves a direct gate swap with y. The three candidates have approximately equal utilities to the actors. Suppose that agent y has exploited in the past and needs to be penalized, and that the penalty is relatively large. Agents w, x and z now add the penalty to their ask prices. Note that the penalty is equal for all three candidates, so the resulting ask prices lie close together. Suppose that candidate 2 with actor w is the cheapest. In this case, it is attractive for either agent x or z to drop its price slightly so that it will be just below that of candidate 2. This will make its candidate the cheapest and thus win the auction instead of candidate 2. Winning the auction is very attractive, since the ask price is still well above the cost price, and the actor in the winning candidate thus makes an attractive profit.

We call this behaviour *forsaking*, purposely deviating from the penalty pricing rule in order to win the auction and make a profit. If agents forsake, two things can happen. First, the exploiters are possibly not fully penalized anymore. Secondly, if multiple agents forsake, they will *compete* against each other. Each forsaker will try to outbid the others by lowering its price. In this way the ask price can go down significantly and the effect of the penalty pricing rule can to a great extent be nullified. How strong the effect of forsaking is depends on the number of forsakers and how often the opportunity occurs. The more forsakers there are, the more will they compete and the less exploiters are penalized. Thus, the more forsakers there are, the more the performance of exploiters will improve. Also, the more forsakers there are, the closer ask prices will be to the disutilities and the smaller the advantage of forsaking will be.

Let σ^f be the forsaker agent type. An agent type vector $\boldsymbol{\sigma} = \langle\, \sigma_1, \sigma_2, ..., \sigma_k \,\rangle$ now specifies agent type $\sigma_i \in \{\, \sigma^c, \sigma^e, \sigma^f \,\}$ for every agent i. Let $\hat{\mathbf{c}}_\mathbf{r} = \langle \hat{c}_r^1, \hat{c}_r^2, \ldots, \hat{c}_r^m \rangle$ be the vector of candidates in ascending order of minimal price, i.e., the vector $\langle\, 1, 2, \ldots, m \,\rangle$ in which the elements are ordered such that $q^{low}_{r,\hat{c}_r^1} \leq q^{low}_{r,\hat{c}_r^2} \leq \ldots \leq q^{low}_{r,\hat{c}_r^m}$, where

$$q^{low}_{r,c} = \begin{cases} -u_{r,c} & \text{if } \ \sigma_{a_{r,c}} = \sigma^f \\ -u_{r,c} + q^p_{r,c} & \text{otherwise} \end{cases}$$

$q^{low}_{r,c}$ denotes the lowest price to which agent $a_{r,c}$ is willing to go. We assume that a forsaker is willing to go as low as his cost price. A forsaker now lowers the penalty if that will make him win the auction. We assume (Q4) from now on, defined as follows.

$$\text{(Q4)} \quad : \quad \forall i, j, r, c\colon \; q_{i,j,r,c} = -u_{i,j,r,c} + q^p_{i,j,r,c} - q^f_{i,j,r,c} + q^{e\prime}_{i,j,r,c}$$

where $q^p_{i,j,r,c}$ is as defined before, and

$$q^f_{i,j,r,c} = \begin{cases} (-u_{i,j,r,c} + q^p_{i,j,r,c}) - q^{low}_{r,\hat{c}^2_r} + \epsilon & \text{if } \sigma_i = \sigma^f \wedge i = a_{r,c} \wedge j = o_r \wedge i \neq j \wedge \\ & \quad c = \hat{c}^1_r \wedge q^{low}_{r,c} < q^{low}_{r,\hat{c}^2_r} < (-u_{i,j,r,c} + q^p_{i,j,r,c}) \\ 0 & \text{otherwise} \end{cases}$$

$$q^{e\prime}_{i,j,r,c} = \begin{cases} q^{low}_{r,\hat{c}^2_r} - q^{low}_{r,\hat{c}^1_r} - \epsilon & \text{if } \sigma_i = \sigma^e \wedge i = a_{r,c} \wedge j = o_r \wedge i \neq j \wedge \\ & \quad c = \hat{c}^1_r \wedge q^{low}_{r,\hat{c}^1_r} < q^{low}_{r,\hat{c}^2_r} \\ 0 & \text{otherwise} \end{cases}$$

Note that exploiters adapt their strategy to the forsakers. When raising their price, they take into account the lowest possible price that the forsakers will ask, and will stay below that price.

Forsakers undermine the penalty pricing rule. The following proposition shows that the presence of forsakers can fully neutralize the effect of retaliation, giving an exploiter back its advantage of exploiting.

Proposition 5.4 (Inequity under forsaking)**.** *Let π denote $\langle A, u, \mathbf{o} \rangle$ and let $k = |A|$. Let $\boldsymbol{\sigma}$ be a type vector such $\sigma_i = \sigma^e$ for some agent i. Then, under* (Q4), (E3)*, there exist π and β such that $\forall r\colon U^-_r < (r-2) \cdot \beta$.*

Proof. The proof is identical to that of proposition 4.1, given the following changes. $A = \{x, y, z\}$ and $\sigma_x = \sigma^f$. Agent y is still an exploiter and exploits in every uneven round after round 1 for the amount of $-2 \cdot \beta - \epsilon$. In rounds in which y is problem owner, x forsakes by asking its cost price, $-2 \cdot \beta$, just as it did in the proof of proposition 4.1. However, this time it asks its cost price as a result of competition with agent z, which has a minimal price of $-2 \cdot \beta + \epsilon$ in every round in which y is problem owner. Note that z may be either coalition or forsaker.

The result is thus the same. In every even round after round 3, the nett transferred utility of x changes with $2 \cdot \beta$. In every uneven round after round 3, the nett transferred utility of agent x changes with ϵ. This gives, for $r > 3$, $U^{nett}_{x,r} = \lfloor \frac{r}{2} \rfloor \cdot 2 \cdot \beta + \lfloor \frac{r-1}{2} \rfloor \cdot 2\epsilon$, which is less than $(r-2) \cdot \beta$. □

It might seem a bit unexpected that the forsaker in the proof above ends up with such a low nett transferred utility. Note however, that the forsaker would not have been better off if he had not forsaken. Also, a more realistic example would have involved three agents; a coalition agent that is repeatedly exploited, an exploiter and a forsaker, that repeatedly prevents the exploiter from being penalized. However, as the example scenario with two agents is already rather complex, we chose not to make it any bigger.

Theorem 5.5 (Dominance of forsaking strategy over coalition strategy)**.** *Let $\boldsymbol{\sigma}$ be a type vector such that $\sigma_i = \sigma^c$ for some agent i, and let $\boldsymbol{\sigma'}$ be a type vector such that $\sigma'_i = \sigma^f$ and $\forall j\colon j \neq i \rightarrow \sigma_j = \sigma'_j$. Then, under* (Q4), (E3)*, for all π, r, it holds that $U^p_{i,r}(\langle \pi, \boldsymbol{\sigma}, -\infty \rangle) \leq U^p_{i,r}(\langle \pi, \boldsymbol{\sigma'}, -\infty \rangle)$.*

Proof. Let $\zeta = \langle \pi, \boldsymbol{\sigma}, -\infty \rangle$ and $\zeta' = \langle \pi, \boldsymbol{\sigma'}, -\infty \rangle$. We write $U_{i,r}$, $Q^{in}_{i,r}$, e_r, etc. to denote $U_{i,r}(\zeta)$, $Q^{in}_{i,r}(\zeta)$, $e_r(\zeta)$, etc., and write $U'_{i,r}$, $Q^{in\prime}_{i,r}$, e'_r, etc. to denote $U_{i,r}(\zeta')$, $Q^{in}_{i,r}(\zeta')$, $e_r(\zeta')$, etc..

For $r = 0$ the property obviously holds. Suppose $U^p_{i,r-1} \leq U^{p\,\prime}_{i,r-1}$ (IH). We prove that in round r, the perceived utility of the forsaker will grow equal or more than that of the coalition agent. We distinguish three cases.

(case 1) $o_r = i$. Since i is not an exploiter and other agents have the same strategy in ζ as in ζ', other agents ask the same price to i, i.e., $\forall c$: $q_{A,i,r,c} = q'_{A,i,r,c}$. Also, $\forall c$: $q_{i,i,r,c} = q'_{i,i,r,c}$. Thus, we have $\forall c$: $q_{A,i,r,c} = q'_{A,i,r,c}$. This results in the same candidate being elected in ζ and ζ', i.e., $e_r = e'_r$. From this follows $u_{i,A,r,e_r} = u'_{i,A,r,e'_r}$. Furthermore, since i is problem owner, we have $q_{i,A,r,e_r} = q'_{i,A,r,e'_r} = 0$. We may now derive

$$\begin{aligned}
U^p_{i,r} &= U_{i,r} + b_{i,r} \\
&= U_{i,r-1} + u_{i,A,r,e_r} + b_{i,r-1} + q_{i,A,r,e_r} - q_{A,i,r,e_r} \\
&= U_{i,r-1} + u_{i,A,r,e_r} + b_{i,r-1} - q_{A,i,r,e_r} \\
&\leq U'_{i,r-1} + u_{i,A,r,e_r} + b'_{i,r-1} - q_{A,x,r,e_r} && \text{(by (IH))} \\
&= U'_{i,r-1} + u'_{i,A,r,e'_r} + b'_{i,r-1} - q'_{A,i,r,e'_r} \\
&= U'_{i,r-1} + u'_{i,A,r,e'_r} + b'_{i,r-1} + q'_{i,A,r,e'_r} - q'_{A,i,r,e'_r} \\
&= U'_{i,r} + b'_{i,r} \\
&= U^{p\,\prime}_{i,r}
\end{aligned}$$

(case 2) $o_r \neq i$ and i does not forsake in round r in ζ'. This implies that either i did not need to lower its price to win the auction in ζ' or that i could not lower its price to win the auction, i.e., $a'_{r,\hat{c}^1_r} \neq i$. In the first case $q_{i,A,r,e_r} = q'_{i,A,r,e'_r}$, $e_r = e'_r$ and $a_{r,e_r} = a'_{r,e'_r} = i$. It follows that $u_{i,A,r,e_r} = u'_{i,A,r,e'_r}$. In the second case some other agent won the auction in both ζ and ζ', giving $e_r = e'_r$, $q_{i,A,r,e_r} = q'_{i,A,r,e'_r} = 0$, $u_{i,A,r,e_r} = u'_{i,A,r,e'_r} = 0$. In both cases, $q_{A,i,r,e_r} = q'_{A,i,r,e'_r} = 0$ since $o_r \neq i$. We may now follow the same derivation as in case 1.

(case 3) $o_r \neq i$ and i forsakes in round r in ζ'. Then, $a_{r,e_r} \neq i$ but $a'_{r,e'_r} = i$. Also, $q_{A,i,r,e_r} = q'_{A,i,r,e'_r} = 0$ since $o_r \neq i$. As $a_{r,e_r} \neq i$, we have $q_{i,A,r,e_r} = 0$ and $u_{i,A,r,e_r} = 0$. As i forsakes in ζ', it holds that $q'_{i,A,r,e'_r} \geq -u'_{i,A,r,e'_r}$, or $q'_{i,A,r,e'_r} + u'_{i,A,r,e'_r} \geq 0$ $(*)$. We may derive

$$\begin{aligned}
U^p_{i,r} &= U_{i,r} + b_{i,r} \\
&= U_{i,r-1} + u_{i,A,r,e_r} + b_{i,r-1} + q_{i,A,r,e_r} - q_{A,i,r,e_r} \\
&= U_{i,r-1} + b_{i,r-1} \\
&= U^p_{i,r-1} \\
&\leq U^{p\,\prime}_{i,r-1} && \text{(by IH)} \\
&= U'_{i,r-1} + b'_{i,r-1} \\
&\leq U'_{i,r-1} + u'_{i,A,r,e'_r} + b'_{i,r-1} + q'_{i,A,r,e'_r} && \text{(by } (*)\text{)} \\
&= U'_{i,r-1} + u'_{i,A,r,e'_r} + b'_{i,r-1} - q'_{A,i,r,e'_r} + q'_{i,A,r,e'_r} \\
&= U'_{i,r} + b'_{i,r} \\
&= U^{p\,\prime}_{i,r}
\end{aligned}$$

□

In section 7, we prove by experiment the expected dominance of forsaking over the coalition strategy. Also, we assess the influence that forsaking has on the effectiveness of the penalty rule. In particular, we test whether, as a result of a decreased effectiveness of penalization, exploitation might become attractive again if there are many forsakers. Also, we measure the level of inequity that will occur in a realistic ATM scenario with exploiters and forsakers.

5.2 Retaliating forsakers

Forsakers enrich themselves at the expense of other agents, and lessen the effect of penalization. The question rises whether forsakers can be penalized just as exploiters, thereby making forsaking unattractive. Unfortunately, from a practical perspective, this is much harder to do.

First, forsaking may go unnoticed. To obtain a significant profit, an exploiter will have to ask a price far above the price others expect, its disutility. Such a difference is easily observable. A forsaker, in order to obtain a large profit, might need to drop its price only slightly. This is less easily observable. A forsaker that has dropped its price slightly might claim that it is correctly enforcing the penalty pricing rule, and that its disutility happens to be slightly lower than the others expect, which is not easily rebutted. While exploitation is more 'obvious', forsaking is more 'subtle'.

Secondly, retaliating forsakers involves more work for the agents than retaliating exploiters. A problem owner that is the victim of exploitation will have a strong incentive to retaliate the wrongdoer. It has been mistreated and has the (human) inclination to want the other to be penalized. Thus, the initiative of retaliation can be left to the problem owner. If an agent forsakes, the problem owner has no incentive to report the forsaking. The forsaker drives the price down, which is advantageous to the problem owner. The other agents do have an incentive to report the forsaking as it undermines the penalizing strategy. The initiative of retaliation thus lies with the other agents. As a result, they should check each others ask prices, which is more work.

Thirdly, penalizing forsakers involves more information sharing, as agents should have knowledge of each others ask prices. Airlines are usually highly secretive about their operational information. We deem it plausible that an airline will be willing to give some information in a 1-on-1 proposal of a deal ("I will wait for five minutes for the amount of 100"). However, to expect an agent to allow all other agents to get this information too and thus get an insight in their operational efficiency, seems implausible.

Fourthly, even if some form of retaliating forsakers is in effect, forsaking can grow gradually worse. As we said, forsaking that involves only small price drops can go unnoticed and can be very tempting. So, all agents might try to secretly forsake with small amounts. In that case they would all ask a price slightly lower than their disutility for a certain comparable action, for instance the gate swap in figure 3. Everyone is asking a slightly lower price now, but since everyone is asking it, it seems to be the 'normal' price. Agents might now want to duck just below this new 'normal' price, which again might go unnoticed since it is only slightly lower. Agents that try to forsake will not accuse others of doing the same because this would point a finger at themselves as they ask comparable prices. In this way there can be a silent, slow, collective decline of the agreement to retaliate forsakers.

The following proposition shows how a small, undetected drop in price can result in large profits for forsakers, whereas exploiters can only obtain small profits by such a price change.

Proposition 5.6 (Retaliating forsakers)**.** *Suppose that price submissions that differ less than some small amount γ of an agents cost price go unnoticed, i.e., agents are able to detect and determine the amount of exploitation and forsaking if $\mid q_{r,c} - q'_{r,c}| > \gamma$ where $q_{r,c}$ is* *the price agent* $a_{r,c}$ *submitted and* $q'_{r,c}$ *is the price it should have submitted, i.e.,* $\forall i,j,r,c:$ $q'_{i,j,r,c} = -u_{i,j,r,c} + q^p_{i,j,r,c}$ *. Suppose that any deviation larger than* γ *is penalized and that agents therefore stay within this limit. Then, after* r *rounds, a forsaker* z *may have realized an advantage in its nett transferred utility of* $(r-4)\cdot\beta$ *as a result of forsaking, while an exploiter* y *can realize an advantage not greater than* $(n-1)\cdot -\gamma$ *as a result of exploiting.*

Proof. The maximal advantage that exploiter y can achieve after r rounds would be achieved by exploiting in $(r-1)$ rounds for an amount of γ, and then spending it all in one round. Thus, $\forall r\colon U^{nett}_{y,r} \leq (r-1)\cdot\gamma$.

For a forsaker however, a scenario can be constructed in which its nett transferred utility grows by approximately $-\beta$ per round. This scenario is depicted in figure 4. In this scenario, the following three things happen subsequently and repeatedly from round three onwards. Between a coalition agent, an exploiter and a forsaker, the exploiter exploits the coalition agent for the maximum amount possible. Then the forsaker makes a large profit by lowering its price and charging less than the full penalty against the exploiter. Then the forsaker buys the maximal amount of utility possible from the coalition agent. We will now give the scenario in more detail. Note that in this scenario, we allow for tasks with a utility equal to zero. As we normally require utilities to be less than zero, we should have used utilities of ϵ instead, but to reduce notational clutter somewhat we allow ourselves this relaxation.

In round one, an exploiter y exploits x for the amount of $-\beta$ while providing a service with $u_{1,e_r} = 0$. This gives $b_{x,1} = \beta$, $b_{y,1} = -\beta$, $U^{nett}_{x,1} = U^{nett}_{y,1} = 0$.

In round two, y is problem owner and both x and z are actor in a candidate. The utilities of the tasks of x and z are both $\beta - \epsilon$. Agent x correctly charges the penalty and submits ask price $-2\beta - \epsilon$. Agent z forsakes and submits -2β, thereby winning the deal. After payments, $b_{z,2} = -2\cdot\beta$ and $U^{nett}_{z,2} = \beta - \epsilon$.

In round three, agent z elects a candidate in which x is actor, with $u_{3,e_3} = 3\beta$. Agent z pays the cost price, -3β. This gives $b_{z,3} = \beta$ and $U^{nett}_{z,3} = -2\cdot\beta - \epsilon$. Agent x has $b_{x,3} = -2\cdot\beta$, $U^{nett}_{x,3} = 3\cdot\beta$.

In round four, five and six, similar steps are repeated, involving the maximal possible amount of exploitation by agent y, the maximal possible profit on forsaking by agent z and the maximal possible nett transferred utility growth for agent z respectively.

As rounds four, five and six are repeated infinitely, we have for $n > 2$, $U^{nett}_{z,n} = (2 + 3\cdot \lfloor\frac{n-3}{3}\rfloor)\cdot -\beta - \epsilon$, which is greater than $(n-4)\cdot -\beta$. □

5.3 Discussion

Throughout this paper we assume that the occurrence of conflicts and repair candidates is independent from the elected candidates. In reality of course, electing a certain candidate will influence future conflicts and candidates. For instance, moving an aircraft from one gate to another eliminates the possibility of a conflict with this aircraft on the first gate, but introduces the possibility of a conflict with this aircraft on the new gate.

We assume that an agent will be capable to estimate to some extent the future consequences of executing a certain candidate. It may use this information in the calculation of its ask price. For instance, if moving to another gate increases the chance on future conflicts, an agent may ask extra money for such an action. In that way, agents can incorporate their expectations of consequences of candidates in the allocation process.

Of course, such considerations make estimating ones true utility harder. Nevertheless, we repeat that agents may use off-line discussion to increase the believability of their price

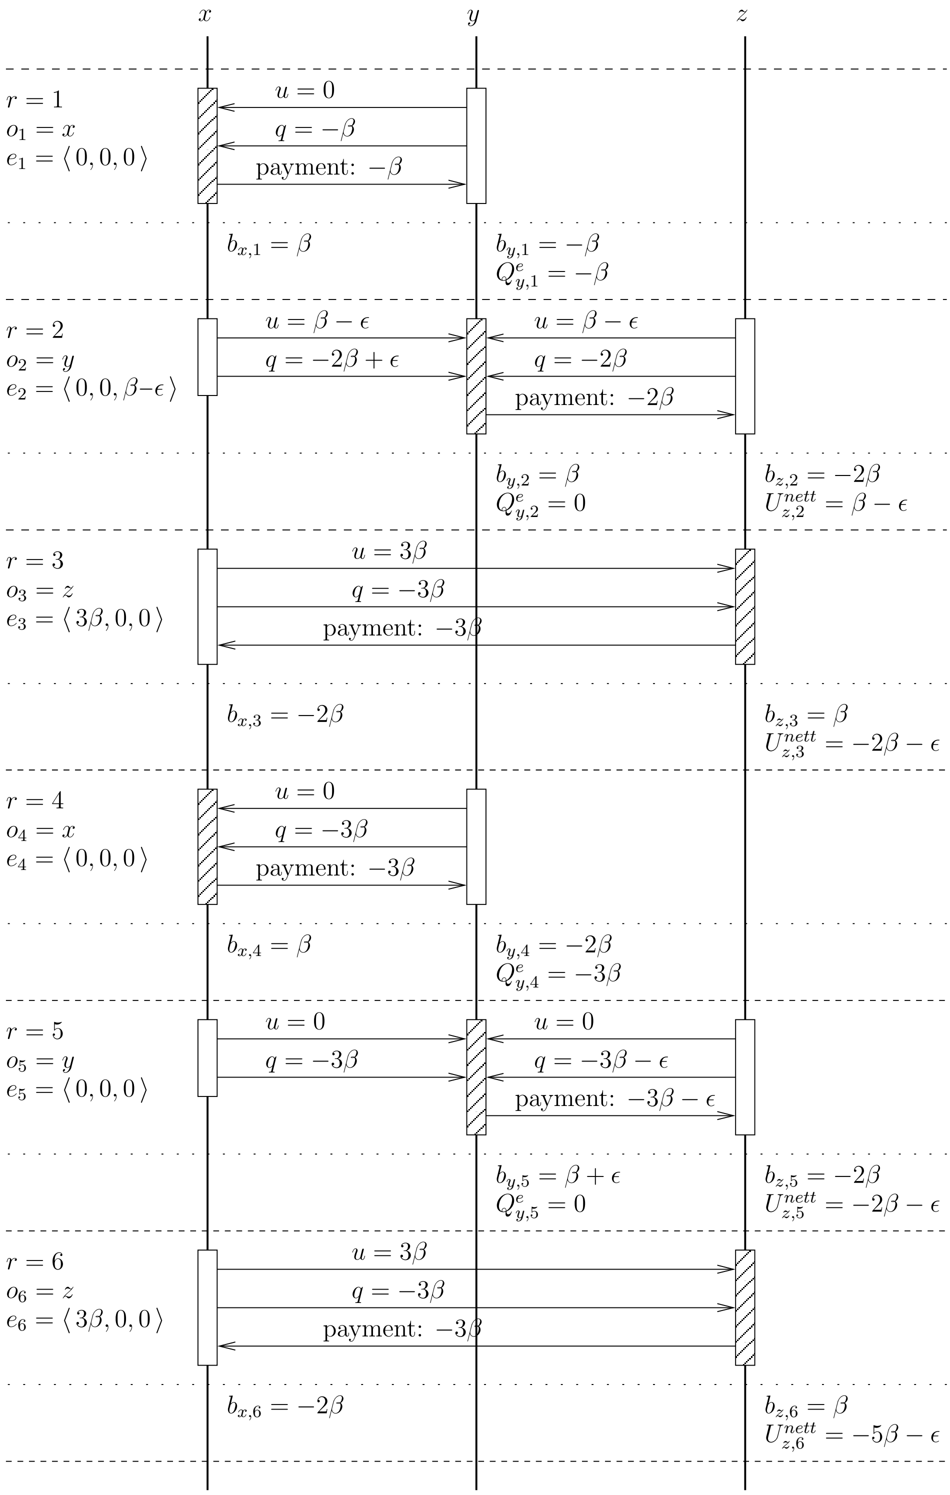

Figure 4: Forsaking with a small price drop. Agent y is exploiter and agent z is forsaker. Shaded bar denotes problem owner.

submission. This may very well include arguments on expected future consequences of candidates.

5.4 Conclusion

Concluding, we can say that the market mechanism with standard currency is not an effective tool to achieve efficient and equitable plan repair in the ATM domain. It would be if agents would truthfully submit cost prices, but selfish agents can be expected to exploit, resulting in an inequitable allocation of repair tasks. The coalition agents may not be able to counteract this via penalty pricing, as it can be attractive to deviate from this rule, thereby undermining it. In the next section we will look at a different way of counteracting exploitation.

6 A new monetary system

The plan repair market discussed so far was based on the use of standard currency. We saw that a market mechanism based on standard currency is unable to robustly counteract exploitation. In this section we solve this problem by moving the penalty system to the currency level. Our proposal comprises a currency system by which exploitation can be counteracted, such that forsaking is not attractive. This currency system is derived from an existing currency system, which we will introduce now.

6.1 WAT System

The system on which our proposal is based is the WAT-system, designed by Eiichi Morino in 2000 [25]. It has since then been used in Japan, both among persons and companies. The very nature of the system makes it hard to assess its scale and number of participants, but it was said to be still in use, and growing, in 2004 [15].

Its most distinctive feature is the fact that it does not need a central bank or administration to keep the books for its users. Instead, users *issue* their own credits by ordering or printing a WAT-ticket and putting their name and signature on it. By signing a ticket, a user vouches for its value, i.e., he promises to exchange a certain amount of goods or services in return when asked. When a user accepts a ticket and wants to spend it in another deal, he adds his name and signature to the ticket. In this way the list of users on a ticket grows as it *circulates*. All these users vouch for the value of the ticket. If the last one should fail to keep its promise, the second to last is liable, and so forth. The longer the list of names on a ticket, the more confidence a user will have that it will yield the promised amount. Finally, when a ticket travels back to its issuer, it is invalidated, which is called *redemption*. Every WAT-ticket goes through the same three stages: issuing, circulation and redemption. An electronic version of the WAT-system called *i*-WAT was developed by Saito in 2003 [18]. The system lets users issue, circulate and redeem tickets electronically. To prevent against fraud, OpenPGP is used and it is required that spending of a credit is approved by the issuer of that credit each time. In the WAT-system, as in most monetary systems, all credits are assumed to have the same value. Although this is highly practical, it is not theoretically sound. A user could very well trust the agents that have signed one credit more than those that have signed another credit, and therefore value the first credit higher. It is understandable that a fixed value is chosen, as it would be infeasible for human users to have to assess the value of every credit individually. For computational agents however, this is not impossible. Their computational power makes it possible for new, complex monetary systems to be used in trade. Our proposal comprises such a currency system, more complex compared to standard currency, but with equity properties that standard currency does not have.

6.2 The Spender-Signed Currency System

In the *spender-signed currency system*, agents may issue their own credits and circulate *foreign* credits. Foreign credits are credits that have been issued by other agents, and have been received in a transaction. Let I_i be the infinitely large set of credits agent i may issue. Let $C_{i,r} = \{ ¢_1, ¢_2, \ldots, ¢_n \}$ be the purse of agent i in round r containing *foreign* credits $¢_1, \ldots, ¢_n$. Let $P_{i,j,r} \subset (I_i \cup C_{i,r-1})$ denote a payment made by agent i to agent j in round r. Thus, a payment may consist of credits issued on the spot or foreign credits that have been obtained before, or both. Let

$$P_{I,J,r} = \bigcup_{i \in I} \bigcup_{j \in J} P_{i,j,r}$$

and

$$P^{out}_{i,r} = \bigcup_{t=0}^{r} P_{i,A,t} \qquad P^{in}_{i,r} = \bigcup_{t=0}^{r} P_{A,i,t}$$

Note that $P_{i,A,t} = P_{i,A \setminus i,t}$ and $P_{A,i,t} = P_{A \setminus i,i,t}$, as an agent never pays any credits to itself. The content of a purse is defined as follows:

$$\forall i, r\colon\ C_{i,r} = \begin{cases} \varnothing & \text{if } r = 0 \\ (C_{i,r-1} \ \setminus \ (P_{i,A,r} \setminus I_i)) \ \cup \ (P_{A,i,r} \setminus I_i) & \text{if } r > 0 \end{cases} \tag{5}$$

Also, we keep track of the credits an agent still has in circulation:

$$\forall i, r\colon\ C^{circ}_{i,r} = \begin{cases} \varnothing & \text{if } r = 0 \\ \Big(C^{circ}_{i,r-1} \ \cup \ (P_{i,A,r} \cap I_i)\Big) \ \setminus \ (P_{A,i,r} \cap I_i) & \text{if } r > 0 \end{cases} \tag{6}$$

Each credit ¢ carries a list of signers $\mathbf{s}_¢ = \langle\, s_{¢,0}, s_{¢,1}, s_{¢,2}, \ldots \rangle$, defined as follows:

$$s_{¢,r} = \begin{cases} i & \text{if } (r = 0 \wedge ¢ \in I_i) \ \vee \ (¢ \in P_{A,i,r}) \\ -1 & \text{otherwise} \end{cases}$$

So, the first signature on a credit is from its issuer. Then, each agent that receives the credit in a payment adds its signature to the credit. If a credit is not spend in a round, the code -1 is used to denote this. For example, a credit ¢ is issued by agent 1 in round 2 and given to agent 4. At that moment, $\mathbf{s}_¢ = \langle\, 1, -1, 4 \,\rangle$. The first value of this sequence shows that ¢ was issued by 1. Round 0 does not exist, in round 1 the credit was not spent and in round 2 it was issued and given to agent 4. After that it stays in the purse of agent 4 in round 3. In round 4 it is given to agent 3, who gives it back to agent 1 in round 5. At that moment, $\mathbf{s}_¢ = \langle\, 1, -1, 4, -1, 3, 1 \,\rangle$. Note that in a practical implementation of the system, each agent will be required to sign a credit before he spends it, hence the name spender-signed currency. We assume here that agents sign a credit directly after having received it, which makes analysing the system slightly easier.

In the spender-signed currency system, each agent i has a reputation $r_i \geq 0$. A reputation may be thought of as a measure for one's trustability, although its theoretical meaning is slightly more complex. A *reputation vector* $\mathbf{r} = \langle r_1, r_2, \ldots, r_k \rangle$ determines for every agent i its reputation r_i.

A credit is valuated by multiplying the reputations of all the agents that signed it up to the current round.

$$v(¢, r) = \prod_{t=0}^{r} r_{s_{¢,t}}$$

where $r_{-1} = 1$. For example, if for a credit ¢, $\mathbf{s}_¢ = \langle 1, -1, 4, -1, 3, 1, -1, -1, -1, ... \rangle$ and $r_1 = 0.9$, $r_3 = 0.8$ and $r_4 = 0.7$, it follows that $v(¢, 4) = 0.9 \cdot 1 \cdot 0.7 \cdot 1 \cdot 0.8 \approx 0.5$.

The value of a set of credits is the sum of the values of the credits in it.

$$v(C, r) = \sum_{¢ \in C} v(¢, r)$$

We assume that values of payments equal ask prices, i.e.,

$$\forall i, j, r\colon\ v(P_{i,j,r}, r{-}1) = q_{j,i,r,e_r} \tag{7}$$

It might seem counterintuitive that the ask price in a certain round must equal the value of a payment set in the previous round. Actually, we require the ask price in a certain round to be equal to the value of a payment set *at the start* of that round. But, as in a round credits may change owner, may be signed and thus may change in value, $v(P, r)$ denotes the value of payment P at the end of round r. To refer to the value of P at the beginning of round r, one needs to use $v(P, r{-}1)$, as is done in (7).

The desirability of an agent's current situation depends on three factors: its cumulative utility, the value of its purse and the expected value of the credits it will have to redeem in the future. The *perceived utility* is defined as follows:

$$U^p_{i,r} = U_{i,r} + v(C_{i,r}, r) - v(C^{circ}_{i,r}, r) \tag{8}$$

A natural equivalent of a bounded balance in the spender-signed currency system is a bound on the value of one's purse minus the value of one's credits in circulation. Hence, we define election rule (E4), an adapted version of (E2), as follows.

$$\text{(E4)} \quad : \quad \forall r\colon\ e_r = \min_c \{\, q_{r,c} \mid v(C_{o_r,r-1}, r{-}1) - v(C^{circ}_{o_r,r-1}, r{-}1) - q^{nett}_{o_r,r,c} \geq \beta \,\}$$

In the spender-signed currency system coalition agents do not charge penalties, but punish exploiters by lowering their reputation, as we will see in the next section. We thus assume (Q2), in which coalition agents ask cost prices and exploiters exploit. For the convenience of the reader, we repeat (Q2).

$$\text{(Q2)} \quad : \quad \forall i, j, r, c\colon\ q_{i,j,r,c} = -u_{i,j,r,c} + q^e_{i,j,r,c}$$

where

$$q^e_{i,j,r,c} = \begin{cases} -(u_{r,\tilde{c}^2_r} - u_{r,\tilde{c}^1_r}) - \epsilon & \text{if } \sigma_i = \sigma^e \wedge i = a_{r,c} \wedge j = o_r \wedge i \neq j \wedge c = \tilde{c}^1_r \\ & \qquad \wedge\, u_{r,\tilde{c}^2_r} < u_{r,\tilde{c}^1_r} \\ 0 & \text{otherwise} \end{cases}$$

Finally, we extend the notion of a scenario with reputations. A scenario $\zeta = \langle \pi, \boldsymbol{\sigma}, \beta, \mathbf{r} \rangle$ is now a tuple consisting of a problem set π, a type vector $\boldsymbol{\sigma}$ and a reputation vector $\mathbf{r}$.

6.2.1 Relation with standard currency

Note that standard currency (with bounded balances) is a specialization of the spender-signed currency system. If we choose $\mathbf{r} = \langle 1, 1, 1, \ldots, 1 \rangle$ and $b_{i,r} = |C_{i,r}| - |C^{circ}_{i,r}|$ $(*)$, all balance-related properties of the standard currency system with exploiters, being (3) (height of balance), (4) (perceived utility) and (E2) (election rule), can be derived. First, we derive that the value of a credit is equal to 1 and that the value of a set of credits is equal to its size.

$$\begin{aligned} v(¢, r) &= \prod_{t=0}^{r} r_{s_{¢,t}} = \prod_{t=0}^{r} 1 = 1 \\ v(C, r) &= \sum_{¢ \in C} v(¢, r) = \sum_{¢ \in C} 1 = |C| \end{aligned} \tag{9}$$

Then we derive (3), (4) and (E2) as follows.

$$\begin{aligned} b_{i,r} &= |C_{i,r}| - |C^{circ}_{i,r}| && (*) \\ &= |\,(C_{i,r-1} \setminus (P_{i,A,r} \setminus I_i)) \cup (P_{A,i,r} \setminus I_i)| - |\,\big(C^{circ}_{i,r-1} \cup (P_{i,A,r} \cap I_i)\big) \setminus (P_{A,i,r} \cap I_i)| \\ & && \text{(by (5), (6))} \\ &= |\,(C_{i,r-1} \setminus (P_{i,A,r} \setminus I_i))\,| - |\,\big(C^{circ}_{i,r-1} \cup (P_{i,A,r} \cap I_i)\big)\,| + |P_{A,i,r}| \\ &= |C_{i,r-1}| - |C^{circ}_{i,r-1}| - |P_{i,A,r}| + |P_{A,i,r}| \\ &= b_{i,r-1} - |P_{i,A,r}| + |P_{A,i,r}| && \text{(by } (*)) \\ &= b_{i,r-1} - v(P_{i,A,r}, r{-}1) + v(P_{A,i,r}, r{-}1) && \text{(by (9))} \\ &= b_{i,r-1} - q_{i,A,r,e_r} + q_{A,i,r,e_r} && \text{(by (7))} \\ &= b_{i,r-1} + q^{nett}_{i,r,e_r} && (3) \end{aligned}$$

$$\begin{aligned} U^p_{i,r} &= U_{i,r} + v(C_{i,r}, r) - v(C^{circ}_{i,r}, r) && (8) \\ &= U_{i,r} + |C_{i,r}| - |C^{circ}_{i,r}| && \text{(by (9))} \\ &= U_{i,r} + b_{i,r} && \text{(by } (*))(4) \end{aligned}$$

$$\begin{aligned} e_r &= \min_c \{\, q_{r,c} \mid v(C_{o_r,r-1}, r{-}1) - v(C^{circ}_{o_r,r-1}, r{-}1) - q^{nett}_{o_r,r,c} \geq \beta \,\} && \text{(E4)} \\ &= \min_c \{\, q_{r,c} \mid |C_{o_r,r-1}| - |C^{circ}_{o_r,r-1}| - q^{nett}_{o_r,r,c} \geq \beta \,\} && \text{(by (9))} \\ &= \min_c \{\, q_{r,c} \mid b_{o_r,r-1} - q^{nett}_{o_r,r,c} \geq \beta \,\} && \text{(by } (*))\text{(E2)} \end{aligned}$$

6.3 Example: only circulation

The most important effect of the spender-signed currency system occurs when money is circulated. We isolate this effect in this section by looking at an agent that only circulates. Note that in a practical implementation of the spender-signed currency system, most agents will probably issue credits. It is however possible that an agent never issues credits.

For an agent i that only circulates credits, i.e., never issued credits, we have that

$$\forall r\colon\ P_{i,A,r} \subseteq C_{i,r-1} \tag{10}$$

As agent i does not issue credits, it will never redeem any credits as well. It follows that

$$\forall r\colon\ (P_{i,A,r}\ \cup\ P_{A,i,r})\ \cap\ I_i = \varnothing \tag{11}$$

From (5) and (11) we may conclude

$$\forall i,r\colon\ C_{i,r} = \begin{cases} \varnothing & \text{if } r = 0 \\ (C_{i,r-1}\ \setminus\ P_{i,A,r})\ \cup\ P_{A,i,r} & \text{otherwise} \end{cases} \tag{12}$$

and $\forall r\colon\ C^{circ}_{i,r} = \varnothing$. As a result, $U^{p}_{i,r}$ reduces to

$$U^{p}_{i,r} = U_{i,r} + v(C_{i,r}, r) \tag{13}$$

For our analysis of the spender-signed currency system, we need to define two new variables. Let $v^*(P^{in}_{i,r})$ denote the sum of the values of the credits received by i up to and including round r, measured for each credit just after it was received by i.

$$v^*(P^{in}_{i,r}) = \sum_{t=1}^{r} v(P_{A,i,t}, t)$$

Let $v^*(P^{out}_{i,r})$ denote the sum of the values of the credits spent by i up to and including round r measured for each credit at the moment it was spent by i.

$$v^*(P^{out}_{i,r}) = \sum_{t=1}^{r} v(P_{i,A,t}, t{-}1)$$

The following invariants are easily checked

$$\forall i,r\colon\ v^*(P^{in}_{i,r}) = v^*(P^{in}_{i,r-1}) + v(P_{A,i,r}, r) \tag{14}$$

$$\forall i,r\colon\ v^*(P^{out}_{i,r}) = v^*(P^{out}_{i,r-1}) + v(P_{i,A,r}, r{-}1) \tag{15}$$

We will determine the relation between the reputation of an agent that only circulates and its perceived utility, as well as give a condition that the reputations must satisfy for exploitation to be unattractive. First, we need to prove the following four propositions.

Lemma 6.1. $v(C_{i,r}, r) = v^*(P^{in}_{i,r}) - v^*(P^{out}_{i,r})$

Proof. For $r = 0$ the proposition clearly holds. Suppose

$$v(C_{i,r-1}, r{-}1) = v^*(P^{in}_{i,r-1}) - v^*(P^{out}_{i,r-1}) \tag{IH}$$

Dividing $C_{i,r-1}$ in two parts, a part that is still in possession of i in round r and a part that he spends in round r, and knowing from (10) that $P_{i,A,r} \subseteq C_{i,r-1}$, we may write

$$C_{i,r-1} = (C_{i,r-1} \setminus P_{i,A,r}) \cup P_{i,A,r} \tag{$*$}$$

It is easy to see that v is distributive. We may write

$$v(C_{i,r-1}, r{-}1) = v((C_{i,r-1} \setminus P_{i,A,r}), r{-}1) + v(P_{i,A,r}, r{-}1) \tag{$**$}$$

Filling in $(*)$ in (12), given that $r > 0$, gives

$$C_{i,r} = (C_{i,r-1} \setminus P_{i,A,r}) \cup P_{A,i,r}$$

Seeing that $C_{i,r-1} \setminus\ P_{i,A,r}$ and $P_{A,i,r}$ are disjoint, we may write

$$v(C_{i,r}, r) = v(C_{i,r-1} \setminus P_{i,A,r}, r) + v(P_{A,i,r}, r) \tag{†}$$

The value of the part that agent i already had in possession did not change. Therefore,

$$v(C_{i,r-1} \setminus P_{i,A,r}, r) = v(C_{i,r-1} \setminus P_{i,A,r}, r{-}1) \tag{‡}$$

We can now rewrite as follows

$$\begin{aligned} v(C_{i,r}, r) &= v(C_{i,r-1} \setminus P_{i,A,r}, r{-}1) + v(P_{A,i,r}, r)) && \text{(by (†),(‡))} \\ &= v(C_{i,r-1}, r{-}1) - v(P_{i,A,r}, r{-}1) + v(P_{A,i,r}, r) && \text{(by (10) and (**))} \\ &= v^*(P^{in}_{i,r-1}) - v^*(P^{out}_{i,r-1}) - v(P_{i,A,r}, r{-}1) + v(P_{A,i,r}, r) && \text{(by (IH))} \\ &= v^*(P^{in}_{i,r}) - v^*(P^{out}_{i,r}) && \text{(by (14),(15))} \end{aligned}$$

□

Lemma 6.2. $v^*(P^{out}_{i,r}) = Q^{in}_{i,r}$.

Proof. For $r = 0$ the proposition clearly holds. Suppose

$$v^*(P^{out}_{i,r-1}) = Q^{in}_{i,r-1} \tag{IH}$$

We have

$$\begin{aligned} v^*(P^{out}_{i,r}) &= v^*(P^{out}_{i,r-1}) + v(P_{i,A,r}, r{-}1) && \text{(by (15))} \\ &= Q^{in}_{i,r-1} + q_{A,i,r,e_r} && \text{(by (IH), (7))} \\ &= Q^{in}_{i,r} \end{aligned}$$

□

Lemma 6.3. $v^*(P^{in}_{i,r}) = Q^{out}_{i,r} \cdot r_i$.

Proof. For $r = 0$ the proposition clearly holds. Suppose

$$v^*(P^{in}_{i,r-1}) = Q^{out}_{i,r-1} \cdot r_i \tag{IH}$$

It is easy to see that

$$v(P_{A,i,r}, r) = v(P_{A,i,r}, r{-}1) \cdot r_i \tag{*}$$

We have

$$\begin{aligned} v^*(P^{in}_{i,r}) &= v^*(P^{in}_{i,r-1}) + v(P_{A,i,r}, r) && \text{(by (14))} \\ &= Q^{out}_{i,r-1} \cdot r_i + v(P_{A,i,r}, r{-}1) \cdot r_i && \text{(by (IH), (*))} \\ &= (Q^{out}_{i,r-1} + q_{i,A,r,e_r}) \cdot r_i && \text{(7)} \\ &= Q^{out}_{i,r} \cdot r_i \end{aligned}$$

□

Proposition 6.4. $U^p_{i,r} = U_{i,r} + Q^{out}_{i,r} \cdot r_i - Q^{in}_{i,r}$

Proof.

$$\begin{aligned} U^p_{i,r} &= U_{i,r} + v(C_{i,r}, r) && \text{(by (13))} \\ &= U_{i,r} + v^*(P^{in}_{i,r}) - v^*(P^{out}_{i,r}) && \text{(by prop. 6.1)} \\ &= U_{i,r} + Q^{out}_{i,r} \cdot r_i - Q^{in}_{i,r} && \text{(by props. 6.3 and 6.2)} \end{aligned}$$

□

The following proposition shows that a higher reputation is beneficial to a circulating agent.

Proposition 6.5 (Circulation)**.** *Let π be a problem set, let $\boldsymbol{\sigma}$ be a type vector. Let i be an agent and let $\mathbf{r}$ and $\mathbf{r}'$ be reputation vectors such that $\forall j : j \neq i \rightarrow r'_j = r_j$. Let $\zeta = \langle \pi, \boldsymbol{\sigma}, -\infty, \mathbf{r} \rangle$ and $\zeta' = \langle \pi, \boldsymbol{\sigma}, -\infty, \mathbf{r}' \rangle$. Then, under* (Q2), (E4)*, for an agent i that only circulates, it holds that $\forall r$: $U^p_{i,r}(\zeta') = U^p_{i,r}(\zeta) + (r'_i - r_i) \cdot Q^{out}_{i,r}(\zeta')$.*

Proof. We write $U_{i,r}$, $Q^{in}_{i,r}$, e_r, etc. to denote $U_{i,r}(\zeta)$, $Q^{in}_{i,r}(\zeta)$, $e_r(\zeta)$, etc., and write $U'_{i,r}$, $Q^{in\,\prime}_{i,r}$, e'_r, etc. to denote $U_{i,r}(\zeta')$, $Q^{in}_{i,r}(\zeta')$, $e_r(\zeta')$, etc..

From proposition 6.4 we have

$$\begin{aligned} U^p_{i,r} &= U_{i,r} + Q^{out}_{i,r} \cdot r_i - Q^{in}_{i,r} \\ U^{p\,\prime}_{i,r} &= U'_{i,r} + Q^{out\,\prime}_{i,r} \cdot r'_i - Q^{in\,\prime}_{i,r} \end{aligned}$$

We rewrite the first equation to

$$U_{i,r} = U^p_{i,r} - Q^{out}_{i,r} \cdot r_i + Q^{in}_{i,r} \tag{$*$}$$

As each agent uses the same strategy in ζ and ζ', we have that $Q^{out}_{i,r} = Q^{out\,\prime}_{i,r}$ and $Q^{in}_{i,r} = Q^{in\,\prime}_{i,r}$. As balances are unbounded, the cheapest candidate is chosen each round in both scenario's, and we have $\forall r$: $e_r = e'_r$ and thus $U_{i,r} = U'_{i,r}$. It follows that

$$\begin{aligned} U^{p\,\prime}_{i,r} &= U_{i,r} + Q^{out}_{i,r} \cdot r'_i - Q^{in}_{i,r} \\ &= U^p_{i,r} - Q^{out}_{i,r} \cdot r_i + Q^{out}_{i,r} \cdot r'_i && \text{(by (*))} \\ &= U^p_{i,r} + (r'_i - r_i) \cdot Q^{out}_{i,r} \end{aligned}$$

□

In the spender-signed currency system, agents can punish an exploiter by lowering its reputation. If the reputations of exploiters are lowered enough, exploitation will be disadvantageous and thereby unattractive. The following theorem gives the condition that the reputations of exploiters must satisfy for exploitation to be unattractive.

Theorem 6.6 (Exploitation)**.** *Let π be a problem set and let $\boldsymbol{\sigma}$ and $\boldsymbol{\sigma}'$ be type vectors such that $\sigma_i = \sigma^c$ for some agent i, $\sigma'_i = \sigma^e$ and $j \neq i \rightarrow \sigma'_j = \sigma_j$. Let i be an agent and let $\mathbf{r}$ and $\mathbf{r}'$ be reputation vectors. Let $\zeta = \langle \pi, \boldsymbol{\sigma}, -\infty, \mathbf{r} \rangle$ and $\zeta' = \langle \pi, \boldsymbol{\sigma}', -\infty, \mathbf{r}' \rangle$. Let r be a round such that agent i has had a chance to exploit in some round before r and suppose that agent i only circulates money up until round r. Then, under* (Q2), (E4)*, for an agent i that only circulates, if $\mathbf{r}$ and $\mathbf{r}'$ both satisfy the condition*

$$(R1_r) \quad : \quad \forall j\colon \begin{cases} r_j < \frac{-U^{out}_{j,r}}{Q^{out}_{j,r}} & \textit{if } Q^{out}_{j,r} > -U^{out}_{j,r} \\ r_j = \frac{-U^{out}_{j,r}}{Q^{out}_{j,r}} & \textit{otherwise} \end{cases}$$

it holds that $\forall r\colon U^p_{i,r}(\zeta) > U^p_{i,r}(\zeta')$.

Proof. Using (Q2) and the definitions of Q^{out} and U^{out}, it is easy to derive

$$\forall i, r\colon \sigma_i = \sigma^c \rightarrow Q^{out}_{i,r} = -U^{out}_{i,r}$$

$$\forall i, r\colon \sigma_i = \sigma^e \rightarrow Q^{out}_{i,r} \geq -U^{out}_{i,r}$$

Once an exploiter i has exploited, it is easy to see that $Q^{out}_{i,r} > -U^{out}_{i,r}$. From proposition 6.4 we have

$$\begin{aligned} U^p_{i,r} &= U_{i,r} + Q^{out}_{i,r} \cdot r_i - Q^{in}_{i,r} \\ U^{p\,\prime}_{i,r} &= U'_{i,r} + Q^{out\,\prime}_{i,r} \cdot r'_i - Q^{in\,\prime}_{i,r} \end{aligned}$$

which we may rewrite to

$$\begin{aligned} U^p_{i,r} &= U^{self}_{i,r} + U^{out}_{i,r} + Q^{out}_{i,r} \cdot r_i - Q^{in}_{i,r} \\ U^{p\,\prime}_{i,r} &= U^{self\,\prime}_{i,r} + U^{out\,\prime}_{i,r} + Q^{out\,\prime}_{i,r} \cdot r'_i - Q^{in\,\prime}_{i,r} \end{aligned}$$

Since i is coalition in ζ, we have $Q^{out}_{i,r} = -U^{out}_{i,r}$. As a result, $r_i = 1$. This gives

$$U^p_{i,r} = U^{self}_{i,r} - Q^{in}_{i,r}$$

Since i is exploiter in ζ', we have $Q^{out\,\prime}_{i,r} > -U^{out\,\prime}_{i,r}$. As a result, $r'_i < \frac{-U^{out\,\prime}_{i,r}}{Q^{out\,\prime}_{i,r}}$. As $Q^{out\,\prime}_{i,r} > 0$, it follows that

$$\begin{aligned} U^{p\,\prime}_{i,r} &< U^{self\,\prime}_{i,r} + U^{out\,\prime}_{i,r} + Q^{out\,\prime}_{i,r} \cdot \frac{-U^{out\,\prime}_{i,r}}{Q^{out\,\prime}_{i,r}} - Q^{in\,\prime}_{i,r} \\ &= U^{self\,\prime}_{i,r} - Q^{in\,\prime}_{i,r} \end{aligned}$$

As exploiting does not influence the elected candidate, we have $U^{self}_{i,r} = U^{self\,\prime}_{i,r}$. As the strategy of other agents is similar in ζ and ζ', we have $Q^{in}_{i,r} = Q^{in\,\prime}_{i,r}$. Thus, we arrive at

$$U^{p\,\prime}_{i,r} < U^{self\,\prime}_{i,r} - Q^{in\,\prime}_{i,r} = U^{self}_{i,r} - Q^{in}_{i,r} = U^p_{i,r} \qquad \square$$

Figure 5 shows the outcome of an example scenario after three rounds. Outer arrows denote money and inner arrows denote utility. In the first round, agent x issued 10 credits to pay agent z for a utility of -10. The credits now bear the signature of x and z and are thus worth $r_x \cdot r_z = 1 \cdot 1 = 1$ a piece. In the second round, agent y exploits and provides a utility of -5 while getting paid the 10 x-credits. The signature of y is now added to the credits, which are no worth $r_x \cdot r_z \cdot r_y = 1 \cdot 1 \cdot 0.5 = 0.5$ a piece. In the third round, agent y uses the

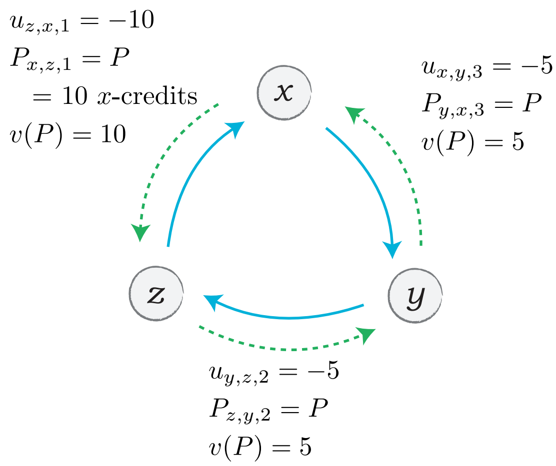

i	x	y	z
r_i	1	0.5	1
$C_{i,0}$	$\varnothing$	$\varnothing$	$\varnothing$
$C_{i,1}$	$\varnothing$	$\varnothing$	P
$C_{i,2}$	$\varnothing$	P	$\varnothing$
$C_{i,3}$	$\varnothing$	$\varnothing$	$\varnothing$
$U^{nett}_{i,0}$	0	0	0
$U^{nett}_{i,1}$	10	0	-10
$U^{nett}_{i,2}$	0	-5	-5
$U^{nett}_{i,3}$	5	0	-5
$U^{p}_{i,3}$	5	0	-5

Figure 5: 10 x-credits are issued, circulated and redeemed. Agent y exploits..

credits it received to pay for a utility of -5, after which the credits are redeemed. It can be seen that agent y is retaliated for its exploitation. Its low reputation resulted in its credits loosing half their value, which led to a final perceived utility of zero.

It can also be seen that the victim of y's exploitation, agent z, has a final perceived utility of -5. In general, the spender-signed currency system does retaliate exploiters, but does not compensate the victims of exploitation.

6.4 Forsaking

An important property of the spender-signed currency system is that forsaking is not an attractive strategy. A forsaker used to win a deal by lowering its bid while still making a profit. In the spender-signed currency system, agents ask cost prices. As a result, lowering one's bid will result in a loss. In other words, there is no common-sense incentive to forsake in the spender-signed currency system.

What forsakers essentially did was slightly deviating from the penalty rule. In the spender-signed currency system, forsakers can do this by lowering reputations of exploiters less than coalition agents do. In this way, they will probably win some deals, as their services appear to be cheaper in the eyes of exploiters than comparable services of coalition agents. But, the credits that they earn in this way have lost value in the eyes of non-forsakers, as the signature of the exploiter is on it. If they spend this money on non-forsakers, it will yield less utility than what they originally provided for it. Thus, forsaking will lead to losses.

We will in section 7 prove by experiment that forsaking is indeed unattractive.

6.5 Example: no circulation

In section 6.3 we saw how under the spender-signed currency system, an agent's reputation led to advantages or disadvantages when the agent circulated money. By setting the reputation of an exploiter to a low value, an exploiter can be penalized and if the losses it makes on circulation are greater than the profits it makes from exploiting, exploitation will be unattractive.

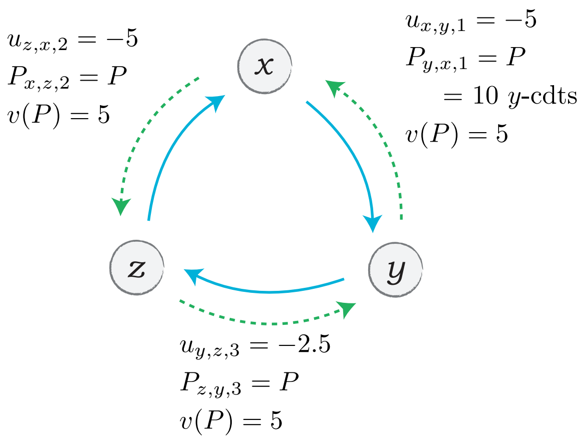

i	x	y	z
r_i	1	0.5	1
$C_{i,0}$	∅	∅	∅
$C_{i,1}$	P	∅	∅
$C_{i,2}$	∅	∅	P
$C_{i,3}$	∅	∅	∅
$U^{nett}_{i,0}$	0	0	0
$U^{nett}_{i,1}$	-5	5	0
$U^{nett}_{i,2}$	0	5	-5
$U^{nett}_{i,3}$	0	2.5	-2.5
$U^{P}_{i,3}$	0	2.5	-2.5

Figure 6: 10 y-credits are issued, circulated and redeemed. Agent y exploits..

However, in the spender-signed currency system as proposed thus far, a low reputation does not negatively affect an agent when it issues and redeems money. To see this, consider the scenario depicted in figure 6. Here, agent y is an exploiter and receives a utility of -5 from by agent x. As $r_y = 0.5$, agent y needs to issue 10 credits. These credits are then circulated to agent z in return for the corresponding amount of utility. Then, in round 3, the credits are redeemed by agent y. We assume that the utility provided by y is -2.5 and that agent y is able to exploit for an amount of 2.5. At the end of round 3, all credits are redeemed so all purses are empty. Agent y has a positive nett transferred utility however, and agent z, the victim of exploitation, a negative nett transferred utility. Thus, agent y was not hindered by its low reputation. For any other reputation of y greater than zero, a similar scenario can be constructed.

In the spender-signed currency system, the reputation of an exploiter, if it is greater than zero, does not have any negative effect if it only issues and redeems credits. One could argue that, in the spender-signed currency system, the only correct valuation of a credit issued by an under-deliverer is also zero. If it were not, the under-deliverer could always make a profit on redemption by providing less utility in return than the worth of that credit. A credit valuation function expressing this principle is the following.

$$v(c, r) = \begin{cases} 0 & \text{if } s_{c,0} < 1 \\ \prod_{t=0}^{r} r_{s_{c,t}} & \text{otherwise} \end{cases} \tag{16}$$

If this valuation function is in effect, an agent that wants to continue to be able to issue credits would have to refrain from under-delivering in order to keep its reputation equal or greater than 1. The disadvantage of such a valuation function is that, as soon as an agent has a reputation less than 1, its issued credits become worthless. If a lot of agents have a reputation less than 1, only few agents can issue money which could stagnate the exchange process by lack of means of payment.

An alternative solution requires *non-uniform reputations*. Before we give this solution, we first need to define the theory for non-uniform reputations.

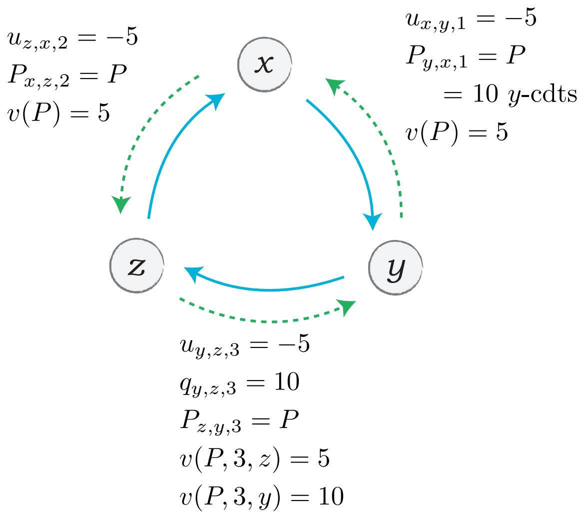

i	x	y	z
$r_{i,i}$	1	1	1
$r_{j\neq i,i}$	1	0.5	1
$C_{i,0}$	$\varnothing$	$\varnothing$	$\varnothing$
$C_{i,1}$	P	$\varnothing$	$\varnothing$
$C_{i,2}$	$\varnothing$	P	$\varnothing$
$C_{i,3}$	$\varnothing$	$\varnothing$	$\varnothing$
$U^{nett}_{i,0}$	0	0	0
$U^{nett}_{i,1}$	-5	5	0
$U^{nett}_{i,2}$	0	5	-5
$U^{nett}_{i,3}$	0	0	0
$U^{p}_{i,3}$	0	0	0

Figure 7: 10 y-credits are issued, circulated and redeemed. Agent y exploits.

6.6 Non-uniform reputations

Non-uniform reputations are the case when each agent i keeps a reputation $r_{i,j} \geq 0$ of each other agent j. Thus, an agent may be considered to have different reputations by different agents. Under non-uniform reputations, the value of a credit ¢ in a round r as perceived by agent i is defined as follows.

$$v(¢, r, i) = \prod_{t=0}^{r} r_{i,s_{¢,t}}$$

A credit will thus be valuated differently by different agents if the reputations that these agents have of the signers differ. Furthermore,

$$v(C, r, i) = \sum_{¢ \in C} v(¢, r, i)$$

The value of a payment set as perceived by the seller should be equal to the ask price.

$$\forall i, j, r\colon\ v(P_{i,j,r}, r{-}1, j) = q_{j,i,r,e_r} \tag{17}$$

Furthermore, we have

$$U^{p}_{i,r} = U_{i,r} + v(C_{i,r}, r, i) - v(C^{circ}_{i,r}, r, i)$$

and we assume (E5), defined as follows

$$\text{(E5)} \quad : \quad \forall r\colon\ e_r = \min_c\{\, q_{r,c} \mid v(C_{o_r,r-1}, r{-}1, o_r) - v(C^{circ}_{o_r,r-1}, r{-}1, o_r) - q^{nett}_{o_r,r,c} \geq \beta \,\}$$

6.6.1 Exploitation on redemption

In section 6.5 we showed how an exploiter could, under the spender-signed currency system, exploit on redemption without being penalized. Under non-uniform reputations, this problem can be solved by requiring agents to have a reputation of themselves equal to 1, i.e.,

$$\forall i\colon\ r_{i,i} = 1 \tag{18}$$

To see the effect of this, consider figure 7, in which a scenario similar to that of figure 6 is depicted. We assume $r_{x,x} = r_{y,x} = r_{z,x} = 1$, $r_{x,z} = r_{y,z} = r_{z,z} = 1$, $r_{x,y} = r_{z,y} = 0.5$ and $r_{y,y} = 1$. The scenario is similar to that of figure 6 until round 3. In that round, agent y asks a price of 10 in return for a utility of -5, thus trying to exploit. He receives from agent z the payment P, which, in his perception, indeed has value $v(P, 3, y) = \sum_{¢\in P} v(¢, 3, y) = \sum_{¢\in P} r_{y,y} \cdot r_{y,x} \cdot r_{y,z} = \sum_{¢\in P} 1 \cdot 1 \cdot 1 = 10$. However, for agent z, the payment P has value $v(P, 3, z) = \sum_{¢\in P} v(¢, 3, z) = \sum_{¢\in P} r_{z,y} \cdot r_{z,x} \cdot r_{z,z} = \sum_{¢\in P} 0.5 \cdot 1 \cdot 1 = 5$. Also for agent x, payment P has value 5. The result is that agent y initially issued the 10 credits in return for a utility of -5, and eventually provided the same utility on redemption. Thus, it has not exploited at all.

7 Experiments

We have in the previous sections shown that in the standard currency system, exploitation cannot be counteracted due to forsakers. The result was an exchange with an undesirable level of inequity. Also we showed that in the spender-signed currency system, exploitation can be counteracted and the resulting exchange is bounded inequitable.

By implementing the plan repair market, both with standard currency and spender-signed currency, we are able to establish four more important results.

Firstly, all of the dominance theorems we gave were under the assumption of unbounded balances. We test by means of experiment whether the dominance results also hold under bounded balances.

Secondly, we assess the effect that bounding the balances has on the efficiency of the resulting exchange.

Thirdly, we assess the level of inequity that would occur if exploitation was not counteracted, or if forsaking was not counteracted.

Fourthly, we implement the spender-signed currency system as a proof of concept. We test whether it performs as expected and whether it has an acceptable running time.

7.1 Implementation of the plan repair market

We implemented the plan repair market as described in sections 2 and 3. We chose to incorporate a phenomenon that is important in current day plan repair: that of differently sized airlines. We gave airlines different sizes such that the biggest airline is ten times as big as the smallest. The size of an airline determines the probability of it being problem owner in a round and actor in a given candidate. One's size has, as we will see, important consequences for one's exploitation power.

In each experiment, we have 10 agents ($k = 10$) and 5000 rounds. Let $size_i$ be the size of agent i. We have chosen sizes $50, 45, 40, \ldots, 5$ for agents $1, 2, 3, \ldots, 10$. The chance of being chosen as problem owner for agent i in round r is equal to $size_i \,/ \sum_j size_j$. We have used colours in graphs to be able to discern different agents. However, the graphs will be almost as informative if printed in black and white. Table 1 shows the sizes and colours of the agents.

Each round, one default candidate is generated, consisting of a task for the owner, and two candidates are generated consisting of tasks for two randomly chosen agents. Note that the problem owner may be actor in the second or third candidate. Thus, there can be more than one default candidate. The chance of being chosen as actor in the second candidate for

i	$size_i$	colour of i	
1	50	black	—
2	45	red	—
3	40	blue	—
4	35	yellow	—
5	30	green	—
6	25	purple	—
7	20	light blue	—
8	15	magenta	—
9	10	gray	—
10	5	orange	—

Table 1: Agent sizes and colours.

an agent i in round r is equal to $size_i \,/\, \sum_j size_j$. The same holds for the third candidate. The first candidate has a utility to the owner that is chosen from a gaussian distribution with mean -50 and variance 5. The other two candidates have utilities for the actor chosen from a gaussian distribution with mean -12.5 and variance 5. So, the second and third candidates are usually more efficient.

Each round, a candidate is chosen according to a certain election rule. No candidate can be chosen that would make the balance (or value of the balance) of an agent go below the bound. We have added an extra rule to the experiments to incorporate a principle that can be expected to be important in reality. Agents that have accumulated a lot of wealth have less need than others to earn extra money, as they are far above the bound. Also, they have typically a low nett transferred utility, which indicates that they have given much more help than received. For both reasons, these agents are exempted from providing help until their balance is restored. We have implemented this rule as follows. In every round r in which there is an agent i for which $b_{i,r} > -2 \cdot \beta$ or $v(C_{i,r}, r) - v(C^{circ}_{i,r}, r) > -2 \cdot \beta$, two things happen. First, agent i is exempted from providing help by bypassing the candidates in which it is actor. Secondly, if agent i is problem owner, it will always choose a non-default candidate if possible, even if a default candidate would be less expensive. In this way the agent uses its money to avoid effort.

7.2 Example scenario's

We start with five example experiments. The first experiment is the *default scenario*, which represents the current situation at airports, where no collaboration takes place.In each round the default candidate is elected and no payments are made. We implemented this by setting β equal to 0, i.e., if π is as described above, the first experiment is the scenario $\langle \pi, \boldsymbol{\sigma^c}, 0 \rangle$ under (Q1) and (E2). The results can be seen in figure 8. Perceived utilities range from -3575 to -28744 after 5000 rounds. Note that in this experiment $\forall i, r\colon U^{out}_{i,r} = U^{in}_{i,r} = b_{i,r} = 0$ and thus $\forall i, r\colon U^{p}_{i,r} = U_{i,r}$. Naturally, inequity is equal to zero in the default scenario.

The second experiment is the *efficient scenario*, which represents the situation where agents fully collaborate by truthfully submitting cost prices while balances are unbounded. Figure 9 shows the result of scenario $\langle \pi, \boldsymbol{\sigma^c}, -\infty \rangle$ under (Q1) and (E2). Perceived utilities range from -1134 to -11098 after 5000 rounds. Thus, perceived utilities are better than in

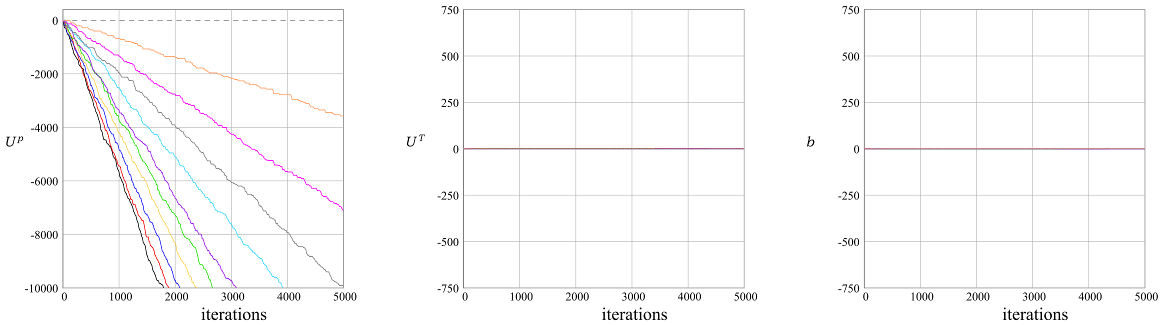

(a) Perceived utilities (b) Nett transferred utilities (c) Balances

Figure 8: Default scenario: $\beta = 0$.

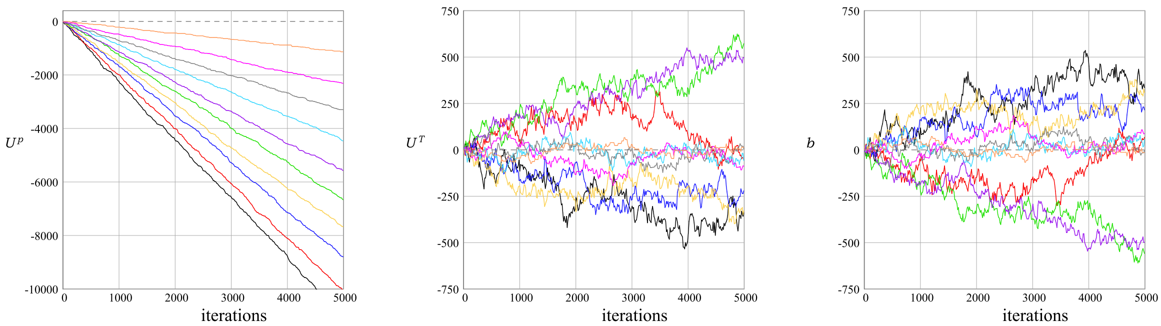

(a) Perceived utilities (b) Nett transferred utilities (c) Balances

Figure 9: Efficient scenario: $\beta = -\infty$.

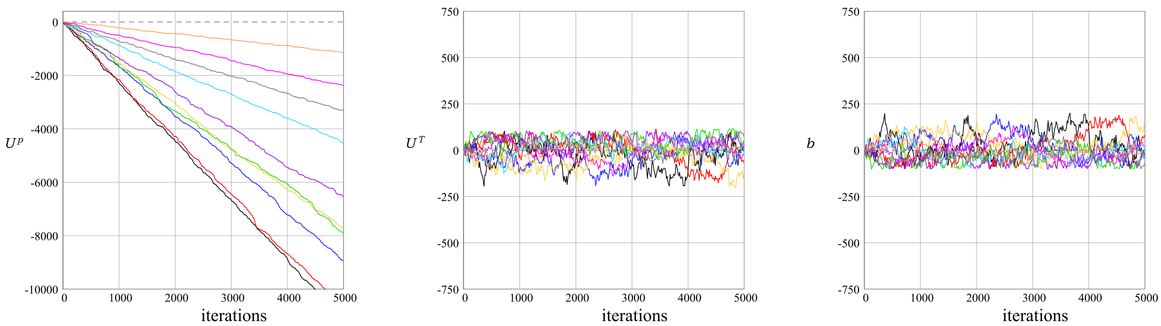

(a) Perceived utilities (b) Nett transferred utilities (c) Balances

Figure 10: Collaborative scenario: $\beta = -100$.

the previous scenario. The invariant of proposition 3.3, $\forall i\colon\ b_{i,r} = U^{nett}_{i,r}$, can be recognized in figures 9b and 9c. The lowest point that U^- reached was approximately -500. Note that in this scenario, U^- is not bounded and may thus become arbitrarily low.

The third experiment is the *collaborative scenario*, in which agents truthfully submit cost prices and use bounded balances. The adjective 'collaborative' denotes the fact that agents truthfully submit cost prices, but do care about equity and hence use bounded balances. Figure 10 shows the result of scenario $\langle\, \pi, \boldsymbol{\sigma^c}, -100 \,\rangle$ under (Q1) and (E2). Perceived utilities range from -1134 to -11352 after 5000 rounds. Thus, perceived utilities are considerably better than in the default scenario and almost as good as in the efficient scenario. Note that the effect of the bound was felt mostly by the three agents who in the efficient scenario had a relatively low balance. These agents had to elect default candidates sometimes to keep their balance above the bound. Furthermore, we see that theorem 3.4 holds, i.e., $\forall r\colon\ U^-_r \geq (k-1)\cdot\beta$. In fact, inequity is at a much more desirable level, the lowest level it reached was approximately $2\cdot\beta$.

The fourth experiment is similar to the collaborative scenario, but two agents have decided to exploit, agents 4 and 8. Figure 11 shows scenario $\langle\, \pi, \boldsymbol{\sigma}, -100 \,\rangle$ under (Q2) and (E2), where $\boldsymbol{\sigma} = \langle\, \sigma^c, \sigma^c, \sigma^c, \sigma^e, \sigma^c, \sigma^c, \sigma^c, \sigma^e, \sigma^c, \sigma^c \,\rangle$. It can be seen that the two exploiting agents achieve a significant improvement in both their perceived utilities and their nett transferred utilities, especially agent 4, the larger of the two.

In the fifth experiment, all agents exploit. Figure 12 shows scenario $\langle\, \pi, \boldsymbol{\sigma}, -100 \,\rangle$ under (Q2) and (E2), where $\boldsymbol{\sigma} = \boldsymbol{\sigma^e}$. It can be seen that the nett transferred utilities diverge, leading to steadily increasing inequity. It can also be seen that the larger airlines come out on top. We will say more about this later.

7.3 Influence of the bound

To examine the relation between the bound, efficiency and equity, we set up a series of experiments in which we vary the bound. Figure 13 shows the results of scenario's $\langle\, \pi, \boldsymbol{\sigma^c}, \beta_0 \,\rangle$, $\langle\, \pi, \boldsymbol{\sigma^c}, \beta_1 \,\rangle$, ..., $\langle\, \pi, \boldsymbol{\sigma^c}, \beta_9 \,\rangle$ under (Q1) and (E2), where $\boldsymbol{\beta} = \langle$ 0, -10, -25, -50, -100, -200, -500, -1000, -2500, -999999 $\rangle$. The first column shows the perceived utilities of the agents in the first experiment, the second column those of the second experiment, and so on. It can be seen that decreasing the bound improves the perceived utilities of the agents but increases inequity.

It is up to the mechanism designer or the agents involved to determine the desired trade-off between efficiency and equity. We judge from figure 13 that $\beta = -100$ is a bound that achieves a desirable trade-off in the collaborative scenario. We will therefore use it in the following experiments. Note that in different scenario's, different bounds might be needed to achieve the most desired trade-off. However, for our experiments a constant bound will suffice as we intend to show that our main results hold under balances that are bounded by *a* bound, not necessarily the ideal bound.

7.4 Dominance of strategies and inequity of exchange under standard currency

To prove dominance of strategies we run several series of experiments. Often experiments with different types of agents are done with a certain chosen distribution of agent types. We however chose to run experiments with many different distributions of agent types. These dis-

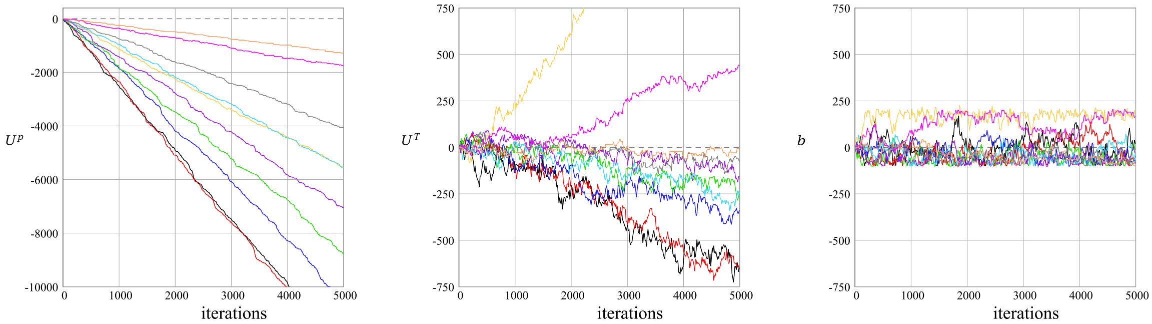

(a) Perceived utilities (b) Nett transferred utilities (c) Balances

Figure 11: Two exploiters.

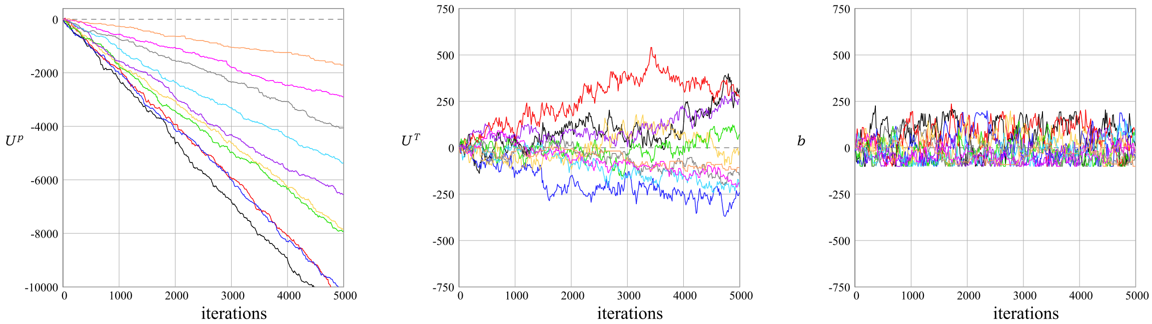

(a) Perceived utilities (b) Nett transferred utilities (c) Balances

Figure 12: All agents exploit.

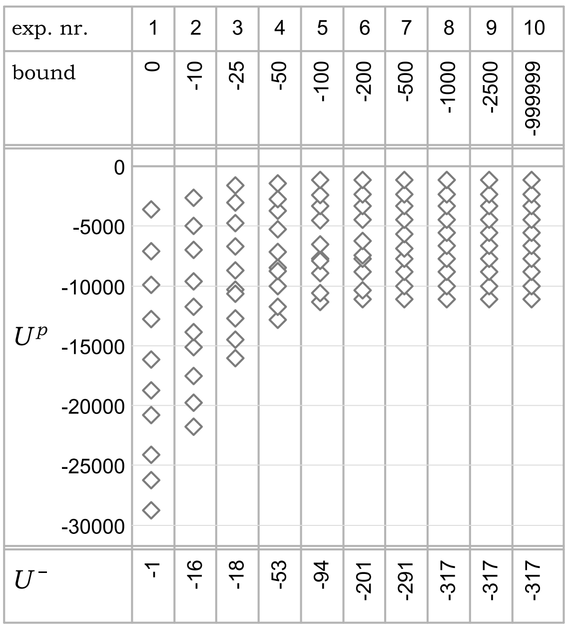

Figure 13: Perceived utilities and equity under a decreasing bound.

tributions are spread uniformly throughout the total distribution space. Extreme cases with only agents of a certain type are covered, and sample distributions between these extremes lie on equal distance from each other. This gives a fairly good sample of the total distribution space, and thereby a good basis to judge the dominance of strategies. Given the problem set, the pricing rule and the bound, if a certain agent type outperforms the other agent types in all of the tested distributions, we have strong evidence that a strategy is dominant.

To correctly compare performance of agents in different scenario's, we use the same problem set in each experiment. This is achieved by seeding the random number generators. Also, we score the agents by their *relative normalized perceived utilities*:

$$\text{score}_{i,r} = \frac{U^p_{i,r} - U^{p\,\prime}_{i,r}}{\text{size}_i}$$

where $U^p_{i,r}$ is the perceived utility of agent i in the current experiment, and $U^{p\,\prime}_{i,r}$ the perceived utility that agent i would have had in the corresponding efficient scenario, i.e., if $\zeta = \langle\, \pi, \boldsymbol{\sigma}, \beta \,\rangle$ is the scenario in the current experiment, $\zeta' = \langle\, \pi, \boldsymbol{\sigma^c}, -\infty \,\rangle$ is the corresponding efficient scenario in which $U^{p\,\prime}_{i,r}$ is measured. Using this definition cancels out the effect of size on an agents score, as well as a priori advantages or disadvantages as a result of the problem set that has been generated (as these a priori advantages occur both in ζ and in ζ').

The inequity measure we have used so far is informative, but has a slight disadvantage. In the collaborative scenario for instance, at a given moment, some agents will have a positive balance and a negative nett transferred utility, and some agents will have it the other way around. As inequity is always measured as the lowest nett transferred utility, inequity will usually be a negative number close to the value of the bound. In the coming experiments, we are most interested in those agents who have a higher or lower nett transferred utility *than they should have* based on their balance. To see which agents have disproportional nett transferred utilities, we use the *adjusted inequity* measure, defined as follows.

$$U_r^{-*}(\mathbf{e}) = \min\{\, U^{nett}_{1,r}(\mathbf{e}) + b_{1,r}, U^{nett}_{2,r}(\mathbf{e}) + b_{2,r}, \ldots, U^{nett}_{k,r}(\mathbf{e}) + b_{k,r} \,\}$$

In the case of non-fungible currency, we use the following definition.

$$U_r^{-*}(\mathbf{e}) = \min\{\, U^{nett}_{1,r}(\mathbf{e}) + v(C_{1,r}, r) - v(C^{circ}_{1,r}, r), \ldots, U^{nett}_{k,r}(\mathbf{e}) + v(C_{k,r}, r) - v(C^{circ}_{k,r}, r) \,\}$$

Thus, the adjusted inequity measure is the inequity measure with remaining balances taken into account.

In our first series of experiments, we let coalition agents compete against exploiters without penalty pricing. We run six experiments, starting from a distribution with only coalition agents, then with increasing number of exploiters, ending with only exploiters. Let

$$\begin{aligned}
\boldsymbol{\sigma^1} &= \boldsymbol{\sigma^c}, \\
\boldsymbol{\sigma^2} &= \langle\, \sigma^e, \sigma^e, \sigma^c, \sigma^c, \sigma^c, \sigma^c, \sigma^c, \sigma^c, \sigma^c, \sigma^c \,\rangle\,, \\
\boldsymbol{\sigma^3} &= \langle\, \sigma^e, \sigma^e, \sigma^e, \sigma^e, \sigma^c, \sigma^c, \sigma^c, \sigma^c, \sigma^c, \sigma^c \,\rangle\,, \\
\boldsymbol{\sigma^4} &= \langle\, \sigma^e, \sigma^e, \sigma^e, \sigma^e, \sigma^e, \sigma^e, \sigma^c, \sigma^c, \sigma^c, \sigma^c \,\rangle\,, \\
\boldsymbol{\sigma^5} &= \langle\, \sigma^e, \sigma^e, \sigma^e, \sigma^e, \sigma^e, \sigma^e, \sigma^e, \sigma^e, \sigma^c, \sigma^c \,\rangle\,, \\
\boldsymbol{\sigma^6} &= \boldsymbol{\sigma^e}
\end{aligned}$$

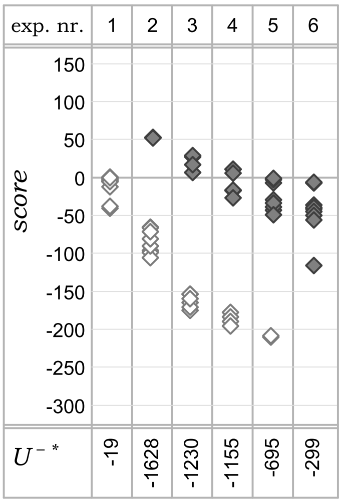

Figure 14: Exploitation.

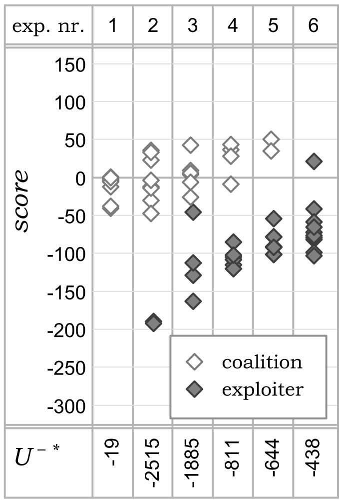

Figure 15: Penalty pricing.

Figure 14 shows the results of experiments $\langle \pi, \boldsymbol{\sigma^1}, -100 \rangle$, $\langle \pi, \boldsymbol{\sigma^2}, -100 \rangle$ through $\langle \pi, \boldsymbol{\sigma^6}, -100 \rangle$ under (E2) and (Q2). It can be seen that in each of the experiments where exploiters participate, they gain a significant advantage over the coalition agents. The experiments confirm the fact that exploitation is dominant over exploiting in the absence of forsakers, when balances are bounded. Examining the agents individually, we observed that bigger agents came out on top. To understand this, we examined the auctions, and found that bigger agents more often were actor in candidates in rounds in which they were problem owner. Thus, they had more often the possibility to solve their own problems, which made them less vulnerable to exploitation. Also, they were more often actor in *both* of the 'cheaper' candidates (candidates two and three) when another agent was problem owner. This made them very strong exploiters, as they could raise their price in both candidates[4]. Thus, the experiment showed that bigger agents are less often exploited and exploit better themselves, which is in line with our reasoning in section 4.1

These differences in exploitation power caused significant differences in nett transferred utilities. When all agents adopt the dominant strategy, i.e., when all agents exploit, adjusted inequity is at an undesirably low level (-299). It can also be seen that the plan repair market with bounded balances is essentially a prisoners dilemma; when all agents follow their optimal, selfish strategy, all agents are worse off than if they would have followed the collaborative strategy.

In our second series of experiments, we let coalition agents use the penalty pricing rule. We choose to make two adjustments to the theory set out in section 5. Firstly, to achieve a non-trivial advantage, the coalition agents structurally overestimate the exploitation done by exploiters slightly. Figure 15 shows the results of experiments $\langle \pi, \boldsymbol{\sigma^1}, -100 \rangle$, $\langle \pi, \boldsymbol{\sigma^2}, -100 \rangle$ through $\langle \pi, \boldsymbol{\sigma^6}, -100 \rangle$ under (E3) and (Q3). It can be seen that under penalty pricing, the coalition strategy is a dominant strategy. When all agents adopt the dominant strategy, i.e., when all agents are coalition, adjusted inequity is at the ideal level (under $\beta = -100$) of -19.

[4]Raising one's price in more than one candidates was not modeled in our theory, as we wanted to keep its complexity to a minimum. We did implement it in the experiments however, as it is realistic behaviour that can be expected from rational, selfish agents.

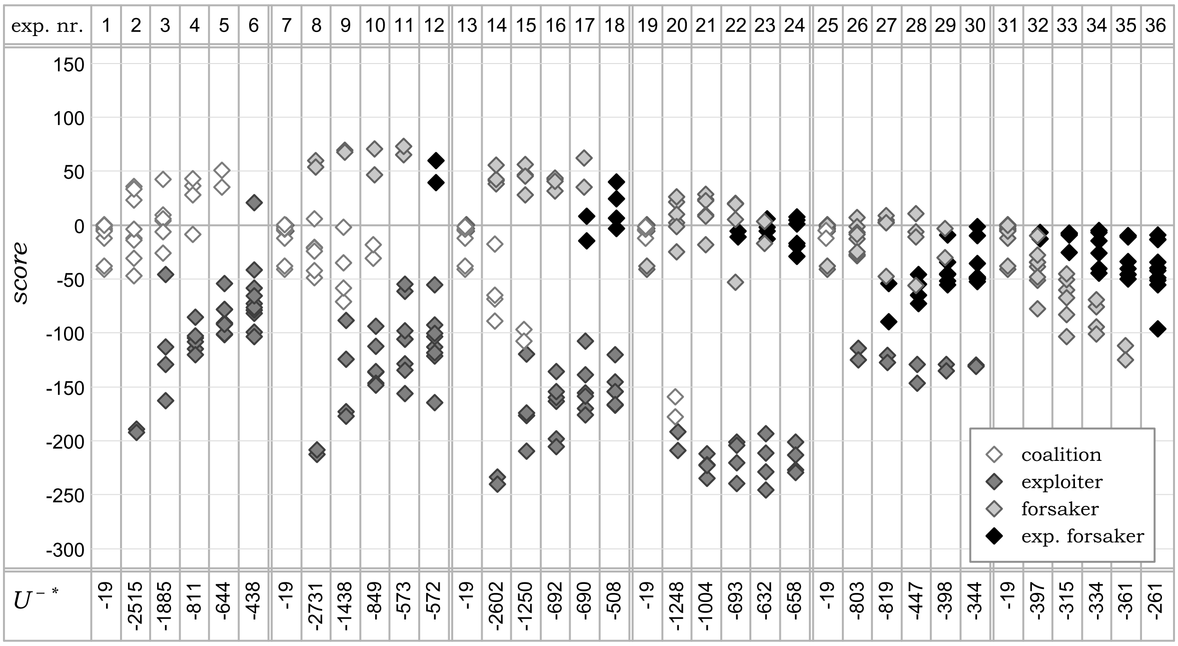

Figure 16: Results of the experiments with forsakers.

In our third series of experiments, we introduce forsaking agents, which exhibit the forsaking behaviour described in section 5.1. We also introduce forsaking exploiters, agents that will exploit in one round and forsake in another. This leads to a larger number of experiments, which can be seen in figure 16. The experiments correspond with scenario's $\{ \langle \pi, \boldsymbol{\sigma}, -100 \rangle \mid \boldsymbol{\sigma} \in \boldsymbol{\Sigma} \}$ under (Q4) and (E3), where $\boldsymbol{\Sigma}$ is the set of all type distributions consisting of m exploiters and n forsakers, where n and m are in $\{ 0, 2, 4, 6, 8, 10 \}$. The first six experiments correspond with the series in figure 15. Subsequent groups of six show increasing numbers of forsakers, until in the last six experiments, all agents are forsaking.

It can be seen that in every experiment, forsaking is a dominant strategy. The last six experiments show that, when all agents forsake, it is dominant to exploit too. Examining the agents individually, we observed that for all agents i except agent 1, $U^p_{i,5000}(\langle \pi, \boldsymbol{\sigma^{ef}}, -100 \rangle) < U^p_{i,5000}(\langle \pi, \boldsymbol{\sigma^{c}}, -100 \rangle)$, where $\boldsymbol{\sigma^{ef}}$ is the type vector denoting that all agents exploit and forsake. Thus, if all agents follow the dominant strategy, all agents except one are worse off. Thus, the plan repair market with standard currency and penalty pricing has the characteristics of a prisoners dilemma. Also, in that case the adjusted inequity is at an undesirably low level (-261).

Concluding, we can say that the experiments confirm our results under bounded balances. Exploitation is dominant over submitting cost prices, and when all agents exploit the inequity of exchange is significant. The penalty pricing rule counteracts exploitation, but forsaking is dominant over the coalition strategy under penalty prizing. When all agents forsake, it is dominant to exploit too. When all agents exploit and forsake, the resulting inequity of exchange is significant.

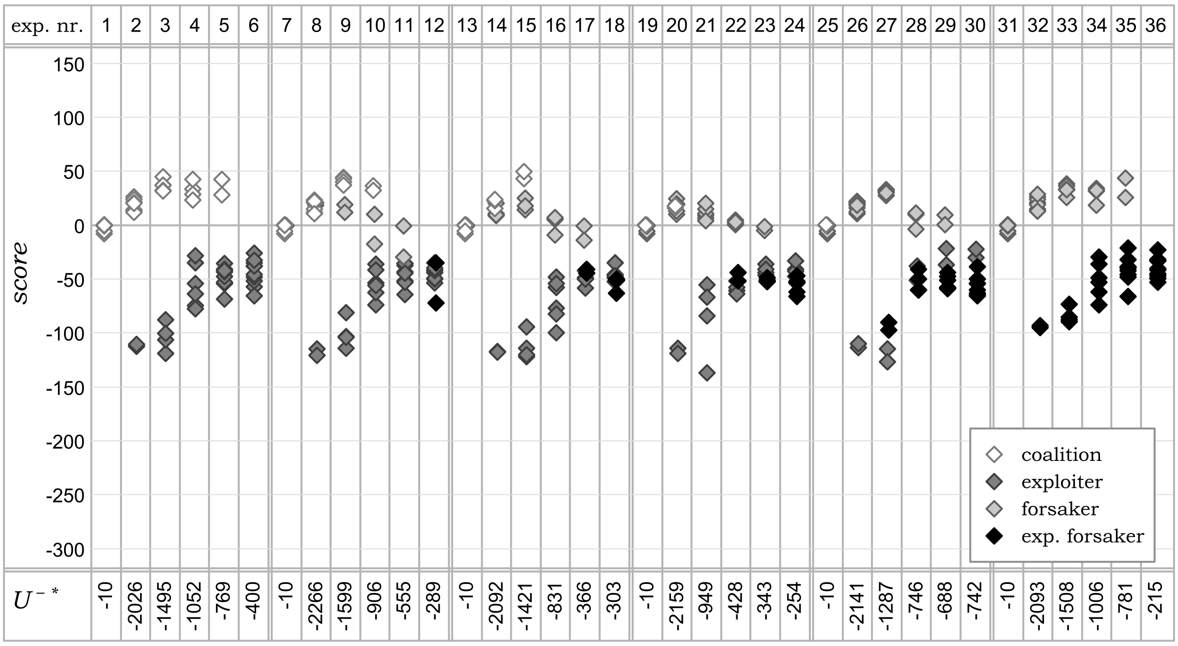

Figure 17: Results of the experiments with spender-signed currency.

7.5 Spender-signed currency system

For our final experiments, we implement the spender-signed currency system. We implement the following dynamic reputation rule.

$$r_{i,r} = \frac{-U_{i,r}^{out}}{Q_{i,r}^{out} + \alpha \cdot (Q_{i,r}^{out} + U_{i,r}^{out})} \tag{19}$$

with $\alpha > 0$. Thus, exploiters are, besides compensated for their exploitation, punished extra proportionally to the amount of exploitation. Most agents use $\alpha = 1.5$. However, there are also forsaking agents, who lower their reputation of forsakers less than the other agents (see also section 6.4). Thereby, they will sometimes win deals. Forsaking agents use $\alpha = 1$.

We test the spender-signed currency system in the same distributions of agent types that we used in the previous series. Figure 17 shows the results of scenario's $\{ \langle \pi, \boldsymbol{\sigma}, -300 \rangle \mid \boldsymbol{\sigma} \in \boldsymbol{\Sigma} \}$ under (Q4) and (E2), where $\boldsymbol{\Sigma}$ is as described earlier. It can be seen that the coalition agents outperform the other agents in every experiment they participate in. This confirms the fact that spender-signed currency is a robust facilitator of efficient and equitable exchange, resistant against forsakers. Reputation rule (19) with $\alpha = 1.5$ is a sufficient reputation rule for its purpose.

This series of experiments took 51.3 seconds to execute on a Xeon 5160, 3.0GHz processor. A total of 4035059 credit transfers took place. As there are 36 experiments with 5000 iterations each, there have been $36 \cdot 5000 = 180000$ rounds. The default candidate was chosen a total of 24654 times, which leaves a total of $180000 - 24654 = 155346$ times a non-default candidate was chosen, i.e., 155346 transactions have taken place. In each transaction, the owner had an average purse size of 225.2. Furthermore, an average of $4035059/155346 = 26.0$ credits were transferred in each transaction. So, to determine an estimated best deal for an owner holding approximately 225 foreign credits, with two potential sellers, with the best deal involving 26.0

credits approximately, costs less than $51.3/155346 = 0.33$ milliseconds. This is more than acceptable for a real-life implementation.

7.6 Conclusion

We saw that the height of the bound influences the trade-off between equity and efficiency, and that a reasonable trade-off between efficiency and equity can be achieved in the collaborative scenario.

The experiments showed that, under standard currency, significant inequity occurs in a scenario with coalition agents, exploiters and forsakers. Forsaking is dominant, and if all agents forsake, exploitation is dominant too. If all agents exploit, bigger agents derive a greater advantage than small agents which leads to significant inequity.

The experiments with spender-signed money showed that the coalition strategy is dominant over exploitation and forsaking and effectively counteracts exploitation if agents use reputation rule (19). The result is that all agents will want to adopt the coalition strategy, which leads to the desired trade-off between efficiency and equity.

8 Discussion and future work

In this paper we chose to use candidates consisting of only a single task to keep complexity to a minimum. We believe that the theory can be extended to candidates that involve multiple tasks for multiple agents. This has interesting consequences for exploitation, coalition and forsaking behaviour. For instance, if there are multiple exploiters actor in the most efficient candidate, they should ‘collaboratively’ exploit if they want to achieve maximal profit. Similar effects occur with penalty pricing and forsaking. We believe that our results will still hold, that is, that exploitation is dominant over asking cost prices, that penalty pricing resolves the exploitation problem, that forsaking is attractive when penalty pricing is in effect and that the spender-signed currency system resolves the forsaking problem.

Another point of future research is the case of a limited instead of an unlimited number of rounds. It might be the case that in practice, the mechanism should be able to ‘restart’, i.e., start over again with cleared balances and counters. This introduces an ‘endgame’ before the mechanism is restarted, possibly with unintended results.

Finally, we want to investigate the efficiency and equity of exchange under the spender-signed currency system when communication and reputations are local instead of global. If agents believe only some of the other agents, and distrust agents they are not familiar with, the system will show a different behaviour. For instance, we expect credits to stay ‘closer’ to the issuer. We want to investigate the effect this has on efficiency and equity.

Also, assuming local communication and trust enables us to answer an important remaining question, namely if the use of reputations introduces new possibilities for strategic behaviour by agents. Among others, we aim to prove the following two hypotheses in future work. Firstly, we hypothesize that an agent will in many cases not derive a significant personal advantage by lowering the reputation it has of another agent. Rather it will find that, as a result, the collaboration between the two will be less. Secondly, we hypothesize that in some situations, lowering the reputation of another agent can be advantageous, but this behaviour can be robustly counteracted just as exploitation can. In general, the spender-signed currency system enables agents to robustly counteract many forms of malicious behaviour.

9 Related work

Currency systems in which agents may issue their own money have been proposed in the literature. Usually these are *issuer-signed* currency systems, in which credits are only identified by their issuer. For instance, in [21] the so-called *lightweight currency protocol* is proposed in which users issue money to pay for usage of resources. The authors do not analyze the economic aspects of their system but expect only a few highly coveted currencies to remain, analogous to the real world. In [16] several token- and trust-based P2P sharing mechanisms are compared and shown to be specializations of an issuer-signed currency system called *stamp trading* [14].

Issuer-signed currency systems suffer from the problem described in section 6.5; malicious agents such as exploiters are not necessarily negatively affected by a declining rate. Moreover, if the rate of a malicious agent has become zero, it can switch to using other currencies and continuing its malicious behaviour. The spender-signed currency system solves this problem by making ones reputation affect circulation.

In research on digital cash, the technique of signing credits is sometimes used to achieve *transferability* of credits, i.e., making it possible to transfer a credits without having to involve the original issuer. Signing credits is then necessary to achieve properties such as anonymity and robustness against fraud [26].

It is often observed that human and corporate behaviour is not purely based on reciprocity but also on commitment [8, 2]. We think that the proposed mechanism is perfectly suited for cooperation based on commitment. A non-uniform reputation value can serve as a commitment value. In that case, the more two agents are committed to each other, the more they find each others bids appealing and the more they will collaborate as a result.

10 Conclusion

We showed that the use of standard currency in the plan repair market can lead to a desirable trade-off between efficiency and equity of exchange among benevolent agents. We also showed that with selfish agents, this trade-off is not guaranteed. When agents exploit, bigger agents have an advantage over small agents, leading to significantly inequitable exchange. We showed that it is not possible for even a large coalition of agents to punish exploiters, as only a few forsakers are enough to undermine the coalition strategy, giving exploiters back their advantage.

We introduced the spender-signed currency system and showed how agents can punish exploiters by adjusting their reputations. Contrary to the standard currency system, forsaking is not attractive as it yields credits that have lost value and therefore lead to a loss. In this system, it is attractive for all agents not to exploit, which gives the desired trade-off between efficiency and equity.

References

[1] J. Adams. Inequity in social exchange. *Advances in Experimental Social Psychology*, 2, 1965.

[2] I. Back and A. Flache. The viability of cooperation based on interpersonal commitment. *Journal of Artificial Societies and Social Simulation*, 9(1), 2006.

[3] M. Ball, R. Hoffman, D. Knorr, J. Wetherly, and M. Wambsganss. Assessing the benefits of collaborative decision making in air traffic management. In *Proceedings of 3rd USA/Europe Air Traffic Management R&D Seminar*, Napoli, Italy, 2000.

[4] M. Ball, R. H. Smith, G. L. Donohue, and K. Hoffman. Auctions for the safe, efficient and equitable allocation of airspace system resources. In P. Cramton, Y. Shoham, and R. Steinberg, editors, *Combinatorial Auctions*, chapter 22, pages 507–538. MIT Press, Cambridge, MA, 2005.

[5] S. J. Brams and A. D. Taylor. *Fair Division: from cake cutting to Dispute Resolution.* Cambridge University Press, 1996.

[6] L. Castellia, P. Debels, and W. Ukovich. Route-charging policies for a central european cross-border upper airspace. *Journal of Air Transport Management*, 11(6):432441, Nov. 2005.

[7] Y. Chevaleyre, P. E. Dunne, U. Endriss, J. Lang, M. Lemaître, N. Maudet, J. Padget, S. Phelps, J. A. Rodríguez-Aguilar, and P. Sousa. Issues in multiagent resource allocation. *Informatica*, 30:3–31, 2006.

[8] M. Emiliani and D. Stec. Aerospace parts suppliers' reaction to online reverse auctions. *Supply Chain Management*, 9(2):139–153, 2004.

[9] T. P. Fan and A. R. Odoni. A practical perspective on airport demand management. *Air Traffic Control Quarterly*, 10(3):285–306, 2002.

[10] J. M. Hoekstra. *Designing for safety: the 'free flight' air traffic management concept.* PhD thesis, Delft University of Technology, Nov. 2001.

[11] G. C. Homans. *Social Behavior: Its Elementary Forms.* Harcourt, Brace & World, New York, 1961.

[12] G. Jonker, J.-J. C. Meyer, and F. Dignum. Efficiency and fairness in air traffic control. In K. Verbeeck, A. Now, B. Manderick, and B. Kuijpers, editors, *Proceedings of the 17th Belgium-Netherlands Conference on Artificial Intelligence (BNAIC '05)*, pages 151–157, 2005.

[13] L. Le, S. Kholfi, G. Donohue, and C.-H. Chen. Proposal for demand management using auction-based arrival slot allocation. In *Proceedings of the 5th USA/Europe Air Traffic Management R&D Seminar*, Budapest, Hungary, June 2003.

[14] R. Levien. Attack resistant trust metrics. draft PhD thesis, 2001.

[15] B. Lietaer. Complementary currencies in japan today: History, originality and relevance. In G. Seyfang and C. Williams, editors, *International Journal of Community Currency Research*, volume 8, 2004. Available at www.uea.ac.uk/env/ijccr.

[16] T. Moreton and A. Twigg. Trading in trust, tokens and stamps. In *Proceedings of the Workshop on Economics of Peer-to-Peer Systems*, Berkeley, June 2003.

[17] J. Sabater and C. Sierra. Reputation and social network analysis in multi-agent systems. In *Proceedings of the First Conference on Autonomous Agents and Multiagent Systems (AAMAS '02)*, pages 475–482, Bologna, Italy, 2002. ACM Press.

[18] K. Saito. Peer-to-peer money: Free currency over the internet. In *Proceedings of the Second International Human.Society@Internet Conference (HSI '03)*, volume 2713 of *LNCS*, pages 404 – 414. Springer-Verlag, Jan. 2003.

[19] J. Schank. Solving airside airport congestion: Why peak runway pricing is not working. *Journal of Air Transport Management*, 11(6):417–425, November 2005.

[20] A. Sen. *Social choice theory*, volume 3, pages 1073–1181. Elsevier, 1986.

[21] D. Turner and K. Ross. A lightweight currency paradigm for the P2P resource market. In *Proceedings of the Seventh International Conference on Electronic Commerce Research*, Dallas, Texas, June 2004.

[22] T. Vossen and M. Ball. Optimization and mediated bartering models for ground delay programs. *Naval Research Logistics*, 53(1):75–90, 2006.

[23] T. Vossen, M. Ball, R. Hoffman, and M. Wambsganss. A general approach to equity in traffic flow management and its application to mitigating exemption bias in ground delay programs. *Air Traffic Control Quarterly*, 11(4):277–292, 2003.

[24] E. Walster, G. Walster, and E. Berscheid. New directions in equity research. *Journal of Personality and Social Psychology*, 25:151–176, 1973.

[25] *WAT-Systems homepage.* www.watsystems.net, 2000.

[26] B. Yang and H. Garcia-Molina. PPay: micropayments for peer-to-peer systems. In S. Jajodia, V. Atluri, and T. Jaeger, editors, *CCS '03: Proceedings of the 10th ACM conference on Computer and communications security*, pages 300–310, New York, USA, 2003. ACM.

[27] H. P. Young. *Equity: In theory and practice.* Princeton University Press, 1994.