

# Achieving Efficient and Equitable Collaboration among Selfish Agents using Spender-Signed Currency \*

## (Short Paper)

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

### ABSTRACT

We study collaboration among selfish agents in the tactical airport planning domain. This can be seen as a social exchange scenario, in which the efforts of performing tasks are the resources that are being exchanged. We investigate conditions under which a market mechanism with the use of standard currency leads to efficient and equitable exchange among benevolent agents. We show that, if some agents are selfish, the mechanism can become inequitable and therefore unacceptable. A straight forward penalty rule is not enough to restore equity, as it is attractive to deviate from such a rule. As a solution, we present an novel currency system, under which malicious agents can be punished, resulting in efficient and equitable exchange.

### Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Multiagent Systems*; K.4.4 [Computers and Society]: Electronic Commerce—*Cybercash, digital cash*

### Keywords

Collaboration, Efficiency, Equity, Alternative Currencies

## 1. INTRODUCTION

Collaborative Decision Making (CDM) is an important trend in Air Traffic Management (ATM) research today. Many innovative proposals in this category however 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 (CDM) techniques successful. Equity is beginning to receive attention in the ATM literature [4], but to the best of our knowledge it has not been addressed in research on CDM yet.

The domain of our research is the phase of planning known as *tactical planning*. This phase of planning is concerned with the sequencing of arriving and departing aircraft and their scheduling

\*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

**Cite as:** Achieving Efficient and Equitable Collaboration among Selfish Agents using Spender-Signed Currency (Short Paper), G. Jonker, F. Dignum and J.-J. Ch. Meyer, *Proc. of 7th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, Padgham, Parkes, Müller and Parsons (eds.), May, 12-16., 2008, Estoril, Portugal, pp. XXX-XXX. Copyright © 2008, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

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*.

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, equity is implicitly required as well; controllers use rules-of-thumb that avoid airlines being affected by problems they did not cause.

In our view, in the new CDM paradigma, plan repair becomes a *social exchange scenario*, in which airlines exchange favours by helping to solve each others problems. 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 [6, 1].

*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. We propose a novel monetary system that can be used, even in the presence of selfish agents, to achieve efficient and equitable exchange.

## 2. EQUITY AND EFFICIENCY

We model the plan repair problem as a *repeated resource allocation* problem. Let  $\mathcal{A} = \{1, 2, \dots, k\}$  be the set of agents. In each round  $r$ , one agent, the *problem owner*  $w_r \in \mathcal{A}$ , has caused a planning conflict. There are a number of ways in which this problem can be solved, called *repair candidates*. Per repair candidate there is one agent that performs a task, thereby solving the conflict. This agent is called the *actor*  $a_{r,j}$  in round  $r$  in candidate  $j$ . Different repair candidates involve different tasks, which have different utilities to the actors. Let  $u_{r,j} \in \mathbb{R}^-$  be the utility that agent  $a_{r,j}$  would incur if repair candidate  $j$  was executed in round  $r$ . For technical convenience, we also use the following notation:

$$u_{i,h,r,j} = \begin{cases} u_{r,j} & \text{if } i = a_{r,j} \wedge h = w_r \\ 0 & \text{otherwise} \end{cases}$$
$$u_{I,H,r,j} = \sum_{i \in I} \sum_{h \in H} u_{i,h,r,j}$$

$u_{i,h,r,j}$  may be read as “the utility that agent  $i$  will incur in aid of agent  $h$  when executing its part in candidate  $j$  in round  $r$ ”. In every round, there is a *default candidate* that assigns a task to the problem

owner, typically with a relatively low utility.

Each round, one of the candidates has to be elected. We write  $e_r$  to denote the candidate that is elected in round  $r$ . An *allocation* is a sequence of elected candidates. Given an allocation  $A = \langle e_1, e_2, \dots, e_n \rangle$ , let

$$U_i(A) = \sum_{r=1}^n u_{i,A,r,e_r} \quad U_{\mathcal{A}}(A) = \sum_{r=1}^n u_{\mathcal{A},A,r,e_r}$$

be the *cumulative utility* of agent  $i$  and the *efficiency* of allocation  $A$  respectively. An allocation  $A = \langle e_1, e_2, \dots, e_n \rangle$  is *efficient* iff for any other allocation  $A'$  it holds that  $U_{\mathcal{A}}(A) \geq U_{\mathcal{A}}(A')$ . Let

$$U_i^{out}(A) = \sum_{r=0}^n u_{i,A \setminus i,r,e_r} \quad U_i^{in}(A) = \sum_{r=0}^n u_{\mathcal{A} \setminus i,i,r,e_r}$$

be the total utility agent  $i$  provided to others and the total utility others provided to  $i$  respectively.

**DEFINITION 1 (EQUITABLE EXCHANGE).** *Allocation  $A$  is equitable iff  $\forall i: U_i^{out}(A) = U_i^{in}(A)$ .*

Let

$$U_i^{nett}(A) = U_i^{out}(A) - U_i^{in}(A)$$

be the *nett transferred utility* of agent  $i$  in allocation  $A$ . As a measure for inequity, we take the often used *egalitarian social welfare*.

**DEFINITION 2 (INEQUITY OF EXCHANGE).** *Inequity of allocation  $A$  is defined as  $U^-(A) = \min\{U_1^{nett}(A), U_2^{nett}(A), \dots, U_k^{nett}(A)\}$*

When it is clear which allocation is referred to, we will use abbreviations  $U_{i,r} = U_i(A_r)$ ,  $U_{i,r}^{out} = U_i^{out}(A_r)$ , etc., where  $A_r$  is the allocation consisting of the first  $r$  entries of allocation  $A = \langle e_1, e_2, \dots, e_n \rangle$  if  $r \leq n$ .

We take the market paradigm as a basis for the plan repair mechanism. In the *plan repair market*, in each round  $r$  the problem owner  $w_r$  opens an auction in which actors may submit ask prices for their tasks. We use  $q_{r,j} \in \mathbb{R}$  to denote the ask price of agent  $a_{r,j}$  in round  $r$  for executing its part in candidate  $j$ . When  $q_{r,j} = -u_{r,j}$  we say that  $q_{r,j}$  is a *cost price*. Again for technical convenience, we also use notation

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

$$q_{I,H,r,j} = \sum_{i \in I} \sum_{h \in H} q_{i,h,r,j}$$

The credits in possession of an agent add up to its *balance*. 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:

$$\forall i, r: B_{i,r} = B_{i,r-1} + q_{i,A,r,e_r} - q_{A,i,r,e_r}$$

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

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 their customers. In our model, let  $B_{min} \in \mathbb{R}^-$  be the lower bound below which the balance of an agent is not allowed to go.

The market mechanism with bounded balances will achieve a 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  $B_{min} = 0$ , only default candidates can be elected. This would be highly inefficient but perfectly equitable. If however  $B_{min} = -\infty$ , 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.

### 3. ENFORCING EQUITY

If agents ask cost prices, the mechanism described will produce the desired outcome. However, selfish agents can *exploit* this mechanism by asking prices above their cost prices if competition allows it. We show in [2] that such behaviour can lead to an ever growing level of inequity.

We are therefore interested to see whether we can adapt the market mechanism in such a way that the coalition strategy, that of truthfully submitting cost prices, is the dominant strategy. This would result in all agents adopting this strategy, which gives the desired trade-off between efficiency and fairness<sup>1</sup>.

A first approach would be to have the coalition agents counteract exploitation by imposing penalties on exploiters. We make the – rather optimistic – assumption that a problem owner that is being exploited is exactly able to observe the level of exploitation. Suppose that all agents share this information among each other and thus know at any moment the amount of money that a certain exploiter  $i$  has earned by exploiting. Let  $B_{i,r}^E \in \mathbb{R}$  be equal to this amount in round  $r$ . Suppose that all agents add this value  $B_{i,r}^E$  to their ask price in any round in which agent  $i$  is problem owner. In addition, they charge an extra, small amount  $q^p \in \mathbb{R}$ . Thus, a coalition agent  $a_{r,j}$  charges  $-u_{r,j}$  but adds  $B_{w_r,r}^E + q^p$  if  $B_{w_r,r}^E > 0$ .

In [2] we analytically prove the dominance of the coalition strategy in the case of unbounded balances and prove by experiment its dominance in the case of bounded balances. The penalty pricing rule manages to make exploitation unattractive.

Unfortunately, the penalty pricing rule is not robust. In some situations, it is attractive for a coalition agent to deviate from it. This is the case if the cost prices of the actors in a certain round lie close together and the problem owner is an exploiter who needs to be penalized. The actors then raise their price with  $B_{w_r,r}^E + q^p$ . Thus, the prices are well above the cost prices of the actors, and lie close together. It can in that case be attractive for an individual agent to lower its price slightly, in order to win the auction. This is beneficial, because the price received is still higher than the cost price. We call this behaviour *forsaking*.

We analytically prove in [2] that forsaking is a dominant strategy over the coalition and exploiting strategy with unbounded balances, and prove by experiment that it is dominant in the case of bounded balances.

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 are not able to counteract this via penalty pricing, as it is attractive to deviate from this rule.

### 4. SPENDER-SIGNED CURRENCY SYSTEM

The problem with the penalty pricing rule is the fact that it is attractive to deviate from it. In this section we propose a currency

<sup>1</sup>The VCG-mechanisms do not apply here, as these mechanisms only maximize efficiency, and not equity, which is of central importance in the ATM domain.

system by which exploiters can be penalized, but in which forsaking is not attractive. Our proposal comprises the formalization and adaptation of an existing currency system, the WAT-system, designed by Eiichi Morino in 2000 [5]. 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 in it. 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 [3].

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.

In the *spender-signed currency system*, agents may issue their own credits and circulate *foreign* credits. Let  $I_i$  be the infinitely large set of credits agent  $i$  may issue. Let  $B_{i,r} = \{c_1, c_2, \dots, c_n\}$  be the purse of agent  $i$  in round  $r$  containing *foreign* credits  $c_1, \dots, c_n$ . Let  $P_{i,h,r} \subset (B_{i,r-1} \cup I_i)$  denote a payment made by agent  $i$  to agent  $h$  in round  $r$ . Let  $P_{I,J,r} = \bigcup_{i \in I} \bigcup_{j \in J} P_{i,j,r}$ . The content of a purse is defined as follows:

$$B_{i,r} = \begin{cases} \emptyset & \text{if } r = 0 \\ (B_{i,r-1} \setminus P_{i,A,r}) \cup (P_{A,i,r} \setminus I_i) & \text{otherwise} \end{cases}$$

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

$$B_{i,r}^- = \begin{cases} \emptyset & \text{if } r = 0 \\ (B_{i,r-1}^- \cup (P_{i,A,r} \cap I_i)) \setminus (P_{A,i,r} \cap I_i) & \text{otherwise} \end{cases}$$

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

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

Each agent  $i$  has a *reputation*  $r_i \geq 0$ . A credit is valued by multiplying the reputations of all the agents that signed it up to the current round.

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

where  $r_{-1} = 1$ . The value of a set of credits is the sum of the values of the credits in it. We assume that values of payments equal ask prices, i.e.,  $\forall i, h, r: v(P_{i,h,r}, r-1) = q_{h,i,r,e,r}$ .

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:

$$U_{i,r}^p = U_{i,r} + v(B_{i,r}, r) - v(B_{i,r}^-, r)$$

In the spender-signed currency system coalition agents do not charge penalties, but punish exploiters by lowering their reputation. We will in the next section give an example of an effective reputation rule.

An important property of the spender-signed currency system is that forsaking is not an attractive strategy. What forsakers essentially did was slightly deviating from the penalty rule. In the spender-signed currency system, forsakers can do a similar thing 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 to non-forsakers, it will yield less than what they originally provided for it. We prove in [2] the dominance of the coalition strategy over both exploiting and forsaking under a natural assumption on the reputation function.

## 5. EXPERIMENTS

By means of experiment, we investigate whether the described behaviour of exploiters and forsakers would have a significant effect on the equity of plan repair. For this we implemented a virtual ATM plan repair scenario. Also, we implemented the spender-signed currency system as a proof of concept, and to test how reputations should be dynamically determined to achieve the desired effect.

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 twice 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.

Each experiment consists of 5000 rounds. There are 10 agents, of which one is randomly chosen to be problem owner each round in such a way that bigger airlines have a greater chance to be chosen. 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. Here also, bigger airlines have a greater chance of being actor. The default candidate has a utility that is randomly chosen from a gaussian distribution with mean  $-40$  and variance  $5$ . The other two candidates have mean  $-10$  and variance  $5$ . So, non-default candidates are usually more efficient.

We implemented two reference experiments. The first reference experiment represents the current situation at airport, i.e., the default candidate is elected in every round and no payments are made. Perceived utilities range from  $-1965$  to  $-27842$  after 5000 rounds.

The second reference experiment represents the situation where agents fully collaborate by truthfully submitting cost prices and electing efficient candidates. Now perceived utilities range from  $-385$  to  $-6774$  after 5000 rounds. Thus, full collaboration would reduce the total burden of plan repair by a factor of  $4.5$  approximately.

To assess the effect of exploitation and forsaking, we implemented exploiters, the penalty pricing rule, and forsakers. 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. This gave us the possibility to empirically observe whether one strategy is dominant or not. If agents with a certain strategy outperform all other agents in all of the tested distributions of agent types, we have strong evidence that the strategy is dominant.

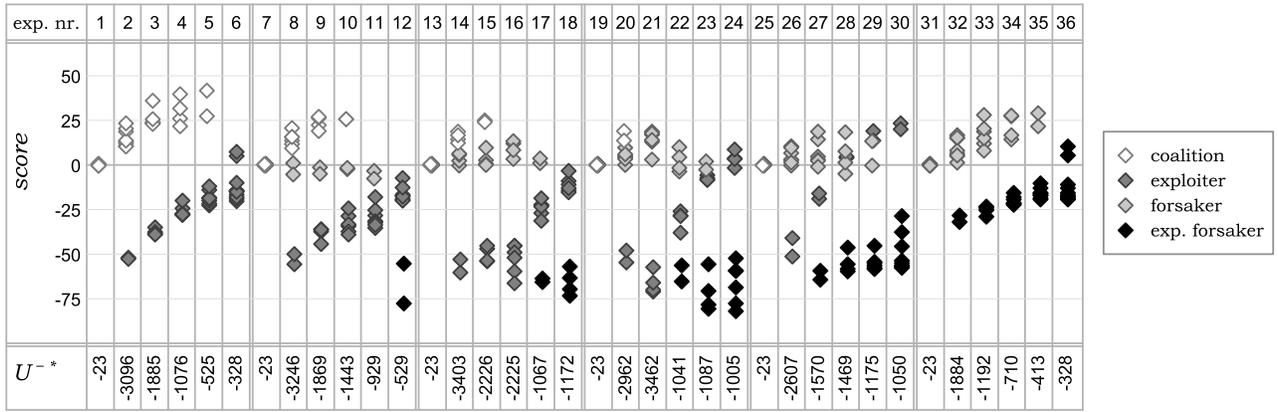


Figure 1: Results of the experiments with spender-signed currency.

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*:

$$\text{score}_{i,r} = \frac{U_{i,r}^p - U_{i,r}^{p'}}{\text{size}_i}$$

where  $U_{i,r}^{p'}$  is equal to the perceived utility of  $i$  that would have occurred if all agents were coalition. Using this definition cancels out the effect of size on an agent's score, as well as a priori advantages or disadvantages as a result of the problem set that has been generated. Thus, in the fully collaborative scenario for instance, all agents would score 0. To measure inequity, we use an adjusted version of definition 2. We take into consideration not only the nett transferred utilities of agents but also their balances.

$$U_r^{-*} = \min \{ U_{1,r}^{\text{nett}} + v(B_{1,r}, r) - v(B_{1,r}^-, r), \dots, U_{k,r}^{\text{nett}} + v(B_{k,r}, r) - v(B_{k,r}^-, r) \}$$

We conducted four series of experiments. Due to space limitations we only show detailed results of the fourth series. In our first series of experiments, we let coalition agents ask cost prices and let exploiters exploit. We ran this test six times, starting from a distribution with only coalition agents, then with increasing number of exploiters, ending with only exploiters. The experiments confirm the fact that exploitation is a dominant strategy.

In our second series of experiments, we let coalition agents use the penalty pricing rule. The experiment confirms the fact that under penalty pricing, the coalition strategy is a dominant strategy.

In our third series of experiments, we introduced forsaking agents, which exhibit the described forsaking behaviour. We ran this test with 36 different agent distributions. The experiment confirms that forsaking is a dominant strategy. Also, the last six experiments show that, when all agents forsake, it is dominant to exploit too. Thus, we can expect all agents to want to exploit and forsake. Importantly, when all agents do this, inequity is at an undesirably low level ( $-513$ ).

In our fourth and final series of experiments, we used the spender-signed currency system as means of payment. Reputations are determined by

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

with  $\alpha > 0$ . Thus, exploiters are, besides compensated for their

exploitation, punished extra proportionally to the amount of exploitation. Most agents use  $\alpha = 2$ . However, there are also forsaking agents, who lower their reputation of forsakers less than the other agents. Thereby, they will sometimes win deals. Forsaking agents use  $\alpha = 1$ . The results are shown in figure 1. It can be seen that the coalition strategy is dominant in every experiment in which coalition agents participate. Thus, under the spender-signed currency system, all agents will want to be coalition. If this is the case, inequity is at a very acceptable level ( $-23$ ).

## 6. CONCLUSION

We showed that the use of standard currency in a market mechanism 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 a great advantage over small agents, which can lead 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.

## 7. REFERENCES

- [1] 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.
- [2] G. Jonker, F. Dignum, and J.-J. Ch. Meyer. Currency systems for efficient and equitable exchange among selfish agents in the air traffic management domain. Technical Report UU-CS-2008-004, Utrecht University, 2008.
- [3] K. Saito. Peer-to-peer money: Free currency over the internet. In *Proc. 2nd Int. Human.Society@Internet Conf.*, volume 2713 of LNCS. Springer, 2003.
- [4] 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.
- [5] WAT-Systems homepage. <http://www.watsystems.net/>, 2000.
- [6] H. P. Young. *Equity: In theory and practice*. Princeton U. P., 1994.