

# Learning from the Metaheuristics: Protocols for Automated Negotiations

Fabian Lang · Andreas Fink

Published online: 22 April 2014  
© Springer Science+Business Media Dordrecht 2014

**Abstract** Nowadays, enterprises are more and more interconnected such that operational planning has to consider the different interests of the involved organizations. This may be a challenging and complex task as it is subject to strategic interactions and incomplete information. Automated negotiation by software agents is a powerful tool which can handle these issues and facilitate intercompany planning. Nevertheless, sophisticated negotiation protocols that govern the rules of the negotiation are needed. In this study, we present and evaluate two configurable protocols for multi-issue negotiations, which are inspired by general heuristic optimization algorithms for centralized problems, so-called metaheuristics. The protocols consist of several policy building blocks; these are evaluated with regard to their impact on the negotiation outcome. The evaluation shows that both protocols can efficiently achieve beneficial solutions—even for complex, nonlinear contract spaces—given the parameterization and the configuration of building blocks are chosen appropriately. Furthermore, we elaborate on requirements for appropriate protocol design and find that both protocols adequately comply with the requirements.

**Keywords** Automated negotiation · Negotiation protocol · Mechanism design · Complex contracts

## 1 Introduction

Negotiation is an important mechanism for the coordination of conflicting interests between multiple agents, be it human or software agents (Jennings et al. 2001).

---

F. Lang (✉) · A. Fink  
Helmut Schmidt University Hamburg, Holstenhofweg 85, 22043 Hamburg, Germany  
e-mail: fabian.lang@hsu-hamburg.de

A. Fink  
e-mail: andreas.fink@hsu-hamburg.de

The complex nexuses of the negotiation domain lead to the need of interdisciplinary research: economists and social scientists are concerned with strategies, tactics, and techniques, whereas computer scientists and information systems researchers are concerned with tools, agents, and platforms (Bichler et al. 2003). Especially the case of automated negotiation gained interest in research in recent years. The inclusion of software agents can facilitate group decision making and make it much more efficient (Lai and Sycara 2008).

Corporate (short-termed) operational planning and decision making is a challenging task. A lot of day-to-day decisions are complex problems (see, e.g., Kallrath 2002). Although necessary information is often available, the companies usually cannot solve their planning problems optimally due to time or computing power constraints. Instead of optimal calculations, they frequently make use of approximation methods and heuristics. Heuristic procedures can be executed in little time and are known to yield satisfying, near-optimal results in many applications. Metaheuristics, general frameworks for designing problem-specific heuristic algorithms, are applicable to a large bandwidth of problems (Blum and Roli 2003). However, nowadays, firms are more and more interconnected—like in supply chains—and face decisions concerning many different entities. The competition of the entities results in malicious strategic interactions and asymmetric as well as incomplete information (Kraus 1996; Raiffa et al. 2002). Consequently, central planning is not feasible anymore. Such intercompany examples are the coordination of production sequences (Fink 2006), supplier-retailer coordination (Rief and Dinther 2010), batch sizing in supply chains (Hombberger 2011), multi-project scheduling with shared resources (Hombberger 2012), or scheduling of shared machines (Lang and Fink 2012c).

Automated negotiation may accomplish such kinds of tasks when central planning is not feasible. Intelligent software agents can find beneficial group decisions on behalf of their principals, the companies or human beings. Negotiation may overcome the challenges resulting from competing negotiators given sophisticated negotiation protocols are in place which can prevent malicious strategic behavior, cope with informational issues, and ensure collective gains from cooperation.

We consider the problem of multilateral multi-issue negotiation with issue interdependencies. Concerning this, we elaborate requirements that, beyond general considerations, also account for the peculiarities of intercompany negotiations. As common procedures mostly do not comply with all those properties, there is a research gap concerning negotiation protocols that meet these requirements and achieve beneficial results. Consequently, the purpose of this study is to present and evaluate negotiation protocols complying with the derived set of prerequisites. Generally, the proposed protocols can be classified as improvement-based single negotiation text procedures, i.e., the contract is iteratively refined and the parties eventually look for Pareto improvements (Raiffa 1982; Vetschera 2013).

The remainder is structured as follows: After this introduction, we define the problem and elaborate on related requirements and performance measures. Following this, we present and discuss related work. Subsequently, two negotiation protocols with a set of optional building blocks are introduced and discussed with regard to the requirements, before we present and elucidate the results of extensive computational experiments. Finally, we conclude the paper and discuss future research needs.

## 2 Problem Description

In the following, we introduce some definitions as well as assumptions to give a structured insight into the negotiation scenario of this study. Furthermore, we develop requirements and objectives for negotiation protocols and present application examples.

### 2.1 Definitions and Assumptions

We consider multi-issue negotiations (see Definition 1) which pose challenging and relevant research field for real-world negotiations (Lai et al. 2006).

**Definition 1** (*Multi-issue Negotiation*) A multi-issue negotiation is a dialog between a set of agents  $\mathcal{J} = \{1, \dots, j, \dots, J\}$  who try to find a consensus about a contract  $c$  that governs the decisions on  $I$  issues by its contract items  $1, \dots, I$  (with  $I > 1$ ).

Multi-issue negotiations are opposed to single issue ones such as the most prominent example, price negotiations. Nevertheless, pricing can be an issue in multi-issue contracts as well—such as in the case of customized products. Depending on the number of agents, negotiations are divided into bilateral (two agents) and multilateral ones (more than two agents).<sup>1</sup>

The contract that is the result of the negotiation assigns decisions to every issue (see Definition 2). Generally, the decision space  $\mathcal{D}_i$  for an issue can be continuous or discrete.

**Definition 2** (*Contract*) A specific contract assigns a decision  $d_i$  from the domain  $\mathcal{D}_i$  to each contract item  $i \in \mathcal{I} : c = \{d_1, \dots, d_i, \dots, d_I\}$  with  $d_i \in \mathcal{D}_i$ .

Referring to the respective decision spaces, the overall contract space comprises all possible contract arrangements (see Definition 3).

**Definition 3** (*Contract Space*) The contract space  $\mathcal{C}$  is the set of all feasible contracts:

$$\mathcal{C} \subseteq \mathcal{D}_1 \times \dots \times \mathcal{D}_i \times \dots \times \mathcal{D}_I.$$

The number of possible contracts in the contract space is  $|\mathcal{C}| \leq \prod_{i \in \mathcal{I}} |\mathcal{D}_i|$ . If the decision spaces have the very same cardinality for every item ( $|\mathcal{D}_i| = |\overline{\mathcal{D}}| \forall i \in \mathcal{I}$ ), the contract space size is given by  $|\mathcal{C}| \leq |\overline{\mathcal{D}}|^I$ .

The following two assumptions are not required for the negotiation protocols as described later, but are mainly introduced for evaluation purposes. The proposed negotiation procedures basically involve pairwise comparisons of contracts by individual agents, thus, the solution approach principally builds on the weaker assumption of ordinal preference relations. However, to assess the overall outcome of a negotiation

<sup>1</sup> There is also a multi-bilateral case with more than two agents, but only bilateral connections (one-to-many).

in an experimental setting, it is useful to be able to compute some measure of social welfare, which leads to the following two supplementary assumptions.

A cardinal utility function maps for an individual agent each element of the contract space to a monetary value (French 1986; see Assumption 1). Consequently, if an agent  $j$  strictly prefers a contract  $c$  to  $c'$  ( $c \succ_j c'$ ),  $U_j(c)$  is larger than  $U_j(c')$ .

**Assumption 1** (*Cardinal Utility Function*) An agent  $j$  has a cardinal utility function  $U_j(c) : \mathcal{C} \rightarrow \mathbb{R}$ .

By assuming that agents' utilities are transferable, which is rather common in game theory (Bergstrom and Varian 1985), respective utility functions lead to cardinal values that can be compared interpersonally and be used to measure a social welfare function that aggregates individual utilities (see Assumption 2). This concept traces back to the seminal work of Neumann and Morgenstern (1944). Transferable utility means that the agents' utility values are measured in terms of a common (usually monetary) commodity and the utility function for the commodity is linear, i.e., all agents have a constant marginal utility for an extra unit. As a result, the agents could transfer their utilities by reallocating this commodity (Kaneko 1976; Bergstrom and Varian 1985; Shoham and Leyton-Brown 2009).

**Assumption 2** (*Transferable Utility*) The utility functions of the agents are transferable, i.e., agents can transfer any given amount of their utility to other agents by transferring a commodity. Consequently, when agent  $j$  transfers the amount  $x$  to agent  $-j$ , the utility of  $j$  is decreased and the utility of  $-j$  is increased by  $x$  units.

Real-world negotiations are often characterized by interdependencies between contract items (Klein et al. 2003), which we consequently assume in our scenario (see Assumption 3). Interdependent items can lead to nonlinear contract spaces—so called complex contracts (Klein et al. 2003). Such contracts are generally characterized by multiple local optima, i.e., the negotiation search space has many “hills” and “valleys”. Those contracts are complex, because, on the one hand, they comprise a very large number of feasible contracts (usually exponential in the size of the problem), and, on the other hand, the local optima exacerbate finding an efficient agreement (Lang and Fink 2012c).

**Assumption 3** (*Interdependencies of Contract Items*) The contract items are interdependent, i.e., the valuation  $V_j(d_\varphi)$  of an agent  $j$  for an item  $\varphi$  can vary depending on the other contract items' decisions  $d_i : i \in \mathcal{I} \setminus \{\varphi\}$ .

In microeconomic theory, this concept of interdependencies is known as complementary and substitute goods (Varian 2010). Agents benefit when complements are combined (e.g., noodles and sauce), i.e., they are superadditive:  $V_j(p \wedge q) > V_j(p \wedge \neg q) + V_j(\neg p \wedge q)$ . In contrast, in the case of substitutes, they lose utility because usually one good is used instead of the other (e.g., noodles and potatoes), i.e., they are subadditive:  $V_j(p \wedge q) < V_j(p \wedge \neg q) + V_j(\neg p \wedge q)$ . This can be extended to higher degrees of combination.

Furthermore, we assume that all preference information is private, i.e., it is just known to the respective agent (see Assumption 4). To prevent strategic disadvantages,

one agent might not want to reveal its information to other agents (Bichler et al. 2003; Fujita et al. 2010a, b; Sandholm 1999). If self-interested agents are free to reveal their preferences and strategies, there may be an incentive to report it untruthfully (Conitzer and Sandholm 2004; Neumann 2007; Vulkan 1999). Although there are mechanisms to elicit the preferences of the agents truthfully like the Vickrey auction (Vickrey 1961), the revelation is not always desirable due to privacy considerations (Bichler et al. 2003; Rothkopf et al. 1990).

**Assumption 4** (*Private Information*) The preferences and the strategy considerations of the agents are private, i.e., they are not known to the other agents.

Commonly, game-theorists and behavioral economists assume rational behavior (see Neumann and Morgenstern 1944; Hurwicz 1945). Software agents are programmed by their human principals which have an interest in designing the agents optimally (Binmore and Vulkan 1999). As software agents can draw on advanced computational methods like statistics, machine learning, or optimization procedures (Bichler et al. 2010), they virtually can act rationally—whereas human rational behavior is rather bounded (Boutillier et al. 1997). Therefore, we assume rational behavior of the agents (see Assumption 5).

**Assumption 5** (*Rational Behavior*) The agents act rationally, i.e., they aim to achieve the *individually* best outcome.

## 2.2 Objectives

In the following, we present some requirements for negotiation protocols as well as performance measures constituting the objectives of the proposed protocols.

### 2.2.1 Requirements

A protocol design has to fulfill certain normative criteria and requirements. Based on the literature (Jennings et al. 2001; Kraus 2001; Lomuscio et al. 2003; Sandholm 1999; Ströbel and Weinhardt 2003; Rosenschein and Zlotkin 1994), we elucidate seven requirements for protocol design that are summarized in Table 1.

A protocol should not be vulnerable to strategic manipulation of the agents, i.e., it should be strategyproof. A mechanism is strategyproof if no agent has an incentive to act untruthfully and, thus, against their real preferences (Satterthwaite 1975). This concept is often referred to as incentive compatibility (Hurwicz 1973). Every protocol demands certain private information from the agents which should be reported truthfully by them (see req. 1). Consequently, the protocol has to provide adequate incentives for this purpose (Neumann 2007, Chap. 2). A further requirement, which is often associated along with incentive compatibility, is individual rationality (see req. 2). Individual rationality means that participation in a mechanism must be worthwhile compared to not participating (Zlotkin and Rosenschein 1996). Besides not entering the negotiation right from the beginning, the agents could also opt out and leave during the negotiation (Kraus et al. 1995).

**Table 1** Requirements for a negotiation protocol

1	Incentive compatibility	A protocol should be incentive compatible, i.e., the agents truthfully reveal certain private information according to the protocol design
2	Individual rationality	The participation in the negotiation must be beneficial, i.e., if an agent obtained a loss from participating, the agent would not participate
3	Behavioral stability	A protocol should provide incentives so that the agents behave stably in a desired manner
4	Guaranteed success	A protocol should find an agreement under all conditions
5	Simplicity	The optimal strategy of the agents should be easy to learn and apply
6	Privacy	The agents should have to reveal preferably little information
7	Scalability	A protocol should be computationally scalable to a large number of issues and agents

The intention of a protocol is to achieve certain goals. From a model-theoretical view, the achievement of these goals is subject to endogenous variables like agent behavior. Thus, the protocol designer has to take behavior into account and rely on the stability of this behavior (Jennings et al. 2001; Sandholm 1999; see req. 3). Negotiations intend to reach an agreement that is at least acceptable for the negotiation parties because the worst agreement is disagreement (Kraus et al. 1995). Accordingly, a protocol should lead to a consensus (see req. 4). To support the agents, a protocol should not exacerbate discovering an agent's optimal strategy, but rather make it easy to figure this out (see req. 5). So, the computational effort should be minimized, often referred to as computational efficiency (Sandholm 1999). In human negotiations, simplicity also includes the understandability of rules and conventions (Neumann 2007, Chap. 2). In automated negotiations, these should be clear; however, in hybrid systems, where software agents negotiate with human beings, adaptations to human behavior may have to be made which complicates finding the optimal strategy (Ockenfels and Roth 2002).

As mentioned beforehand, the agents have private information (see Assumption 4) and are not willing to make them public. There are revelation mechanisms, but they are not always desirable (see previous section). The disclosure of private information such as preferences is especially sensitive in business applications with several competing companies (Sandholm 1999; Fujita et al. 2010a; Lang and Fink 2013). Therefore, we argue that private information should preferably stay private (see req. 6). Furthermore, a negotiation protocol should provide scalability (Jennings et al. 1998). Scalability is often defined in the context of multiple issues, i.e., a protocol should be computationally capable of handling a lot of contract issues (e.g. Collins et al. 2002; Conitzer 2010; Lai and Sycara 2008). Nevertheless, since it is required by many applications and scenarios, scalability regarding the number of agents should also be given (Jennings et al. 1998; Kraus 1997; Sycara 1998; see req. 7).

### 2.2.2 Performance Measures

The most prominent performance measure for negotiation analysis is **Pareto efficiency** (Jennings et al. 2001; Lomuscio et al. 2003). Pareto efficient (also known as Pareto optimal) means that no agent can be better off without leaving another agent worse off. The set of all Pareto efficient contracts  $\mathcal{P}$  is known as the Pareto frontier ( $c^p \in \mathcal{P} \Leftrightarrow \nexists c \in \mathcal{C} \nexists j \in \mathcal{J} \forall -j \in \mathcal{J} \neq j : U_j(c) > U_j(c^p) \wedge U_{-j}(c) \not\leq U_{-j}(c^p)$ ).

Besides Pareto optimality, there are several other concepts that are regarded as desirable, but simultaneously imply a Pareto optimal solution themselves. One is the **social welfare optimum** (Harsanyi 1955). The objective of social welfare (SW) optimization is the maximization of the sum of utility values ( $c^* = \arg \max_c \sum_{j \in \mathcal{J}} U_j(c)$ ). SW requires a cardinal utility function with transferable utility, i.e., the utility difference between two contracts is measurable—which is given by Assumption 2.

In the classical game-theoretical literature, there are also axiomatic solution approaches that are deduced from desirable properties (axioms) and which are Pareto efficient as well. The most famous is the Nash bargaining solution (Nash 1950). Nash concluded that solutions that are connected to the maximization of the product of the utilities (**Nash product**) conform to the requirements of Pareto optimality, symmetry, and independence of equivalent utility representations as well as irrelevant alternatives ( $c^n = \arg \max_c \prod_{j \in \mathcal{J}} U_j(c)$ ).<sup>2</sup> The Nash solution intends to yield a fair agreement whilst maximizing joint efficiency. Kalai and Smorodinsky (KS) extended that concept and substituted the irrelevant alternative axiom with a monotonicity axiom (Kalai and Smorodinsky 1975). However, this bargaining solution is not easily generalizable to more than two negotiators (Roth 1979) without mitigating some axioms (Imai 1983). Because of this and because the KS solution is similar to Nash's, we focus on the Nash product. Another axiomatic solution is the **Egalitarian** proposed by Kalai (1977). The Egalitarian bargaining solution fulfills the axioms of independence of irrelevant alternatives (dropped by KS) as well as the axiom of monotonicity; however, Nash's axiom of invariance of equivalent utility representations is omitted. The Egalitarian solution follows the maximin principle, i.e., the objective is to maximize the utility of the agent worst off ( $c^k = \arg \max_c \min\{U_1(c), \dots, U_J(c)\}$ ).

### 2.3 Illustrative Applications

The presented generic problem model can represent plenty of real-world application scenarios. In this subsection, we present two potential applications for further illustration of the problem domain.

For instance, consider the case of multi-agent scheduling as presented by Lang and Fink (2012c): Each agent represents a company or department and holds a set of jobs that has to be processed on a single machine or parallel machines within a certain time range. The agents use shared resources like a factory, production machines, or container cranes. There might be sequence-dependent setup times between the jobs, i.e., if, e.g., a machine is supposed to construct a car, it has to be reconfigured if

<sup>2</sup> With the origin as disagreement point.



the previous job was constructing a truck. Since every job has a due date such as a planned delivery date, the agents aim to minimize their total tardinesses. The position of a job in the schedule represents a contract item. The items are highly interdependent because changing a job's position leads to altered setup times as well as completion time changes for all subsequent jobs. Typical domains for this kind of problems are production planning or supply chain coordination (Fink 2006).

Another example from the logistics domain is operational planning between cooperating logistic companies as presented similarly by Sandholm (1999). In this scenario, the agents have to allocate tasks between each other. The agents can represent vehicles of the fleet, the companies themselves, or both (the drivers could be freelancers and work on their own). The agents' objective is to minimize their cost, i.e., they intend to reduce their driving time and route length, while trying to accept as profitable tasks as possible. Here, a contract item is an assignment of an operator to a job. The alteration of an assignment can affect the valuation of other jobs as well: As the newly assigned job can be at a another destination, it can lead to a costly detour. Furthermore, the job can, e.g., exceed the load capacity resulting in additional costs. Consequently, this application example consists of interdependent, complex contract spaces as well.

### 3 Related Work

In this section, we present a selection of related work that also deals with automated multi-issue negotiations with complex contract spaces and protocol design. We also refer to Lai et al. (2004) who provide a broad literature overview of older works on multi-issue negotiations.

Ito et al. (2008) present an auction-based protocol incorporating bids for specific contracts, where the agents identify their bids by random sampling and local search. Similarly, Hattori et al. (2007) explore a huge contract space by iterative bids for a contract, where the bids refer to subspaces which decrease in size. Lang and Fink (2012d) draw on combinatorial auctions incorporating collusion between agents to achieve solutions in limited decision time. As argued in Sect. 2.1, auctions are not always desirable due to potential unwanted information revelation (req. 6), the optimal strategy is not implicitly obvious and simple to apply (req. 5), and, in combinatorial domains, scalability is not always given (req. 7). Fujita et al. (2010a) consider explicitly the information revelation issue. They draw on representatives in large-scale negotiations that have a special position in the negotiation. Becoming a representative is connected to an incentive for agents to reveal private information. Nevertheless, in the case of operational planning between firms (see Fink 2006), we argue that the downside of revelation of private information such as capacity utilization may outweigh potential advantages in the negotiations.<sup>3</sup> This may contradict req. 2, as agents might be forced to undesirably participate in a mechanism, and is not in accordance with req. 6, as more information than as little as possible is revealed. Fujita et al. (2010b) propose

---

<sup>3</sup> For instance, the other companies can adjust their advertising activity by using information deduced from rather unimportant operations management actions, i.e., the information can be linked to a more important context outside of the negotiation.



a distributed mediator protocol, in which the contract space is divided and explored in parallel by means of several mediators. However, for general multi-issue decisions with interdependencies, such as for scheduling problems (see Sect. 2.3), there may be no useful decomposition of the contract space available. Furthermore, information is given to the mediators directly which may collude and obtain full information (see req. 6).

Like proposed in this study, there are protocol approaches that draw on concepts from metaheuristics. Similar to simulated annealing (SA), in the seminal work of Klein et al. (2003, 2007), a mediator controls an iterative negotiation with a single alternative proposal per negotiation round. In this protocol, the agents can respond by strong and weak accept or reject signals and obtain a certain number of tokens for winning mixed vote situations. Klein et al. showed for two agents by experiments that truthful responses are superior to exaggerations. The annealing is conducted by the mediator based on the responses of the agents; thus, the prisoner's dilemma can be circumvented. Similarly, Fujita et al. (2014) present a related protocol based on issue grouping and apply a limit on strong votes instead of tokens. However, there are some drawbacks: it is not clear whether the token scheme or the limited votes rule is incentive compatible for multilateral negotiations (req. 1). By using strong and weak responses, there is more revealed information (req. 6). Furthermore, the optimal strategy is not obvious (req. 5) as the agents have to log their actions over time and need intelligent mechanisms to fulfill the protocol conditions in a best possible way. Moreover, the determination of the number of tokens or number of strong votes is not trivial and can change the outcome. Based on the work of Klein et al. (2003), Fink (2006) proposes a Simulated Annealing based protocol, in which the agents have to accept a certain quota over time, e.g., they have to accept one hundred proposals out of one thousand proposal iterations. This protocol considers privacy issues and is easily scalable to more than two agents. By means of quotas, the agents can presumably be forced to behave cooperatively, so the prisoner's dilemma can be overcome. However, the incentive compatibility is hard to analyze due to the complexity of strategic interactions (req. 1) and, again, since the quota has to be reached over time, the optimal strategy is not obvious and hard to execute (req. 5). Moreover, there are also some application-oriented studies: Homberger (2010) present a solution for an uncapacitated lot-sizing problem based on the protocol by Fink (2006) and Lang and Fink (2012b, c) adapt a preliminary version of the SA based protocol, which will be presented later on (MNP-SA; see Sect. 4), to single- and multi-machine scheduling problems.

Furthermore, there is some related work that integrates evolutionary mechanisms into negotiations: Tung and Lin (2005) design a mediation service that draws on the concept of recombination and mutation. In this design, the agents send back some of the proposals to a mediator, which recombines those. More application oriented, Homberger (2012) proposes a population-based  $(\mu, \sigma)$  coordination mechanism, also using recombination and mutation, for multi-project scheduling and Homberger (2011) uses a  $(1, \sigma)$  mechanism for supply chain lot sizing. In those protocols, the agents simply decide on the upcoming population by selecting the offspring contract or population, respectively. The structures of those works are similar to the evolutionary protocol proposed in this study (MNP-GA; see Sect. 5); however, behavioral issues, scalability, or information disclosure are not analyzed. Furthermore, as we are going

to show in the next section, the presented protocol introduces several components that can achieve drastically better outcomes than a simple evolutionary procedure (see Sect. 7).

Concluding, although multi-issue negotiations are widely discussed in the literature, we are not aware of any negotiation protocols that consider the elaborated requirements to the same extent as the presented protocols. The study's protocols partly reuse existing approaches and enhance them, while also introducing new ideas. In the following, we will present the protocols and discuss them with regard to the elaborated requirements.

## 4 A Mediated Negotiation Protocol Based on Simulated Annealing

In this section, we present a negotiation protocol that is inspired by the SA metaheuristic. The protocol is based on the seminal work by [Klein et al. \(2003\)](#) (updated version: [Klein et al. 2007](#)) which was extended by [Fink \(2006\)](#). The differences between the protocols are discussed in Sect. 3.

### 4.1 Simulated Annealing

The SA metaheuristic goes back to an algorithm that was used for the simulation of thermal annealing of substances (see [Metropolis et al. \(1953\)](#)). Based on this algorithm, [Kirkpatrick et al. \(1983\)](#) and [Cerný \(1985\)](#) independently of each other developed an optimization heuristic that Kirkpatrick et al. named *Simulated Annealing*. SA is a **location-based** procedure, i.e., the heuristic searches the neighborhood of a current location in the search space for iterative improvement [Laarhoven and Aarts 1987](#). A heuristic, by definition, cannot guarantee optimality; anyhow, [Granville et al. \(1994\)](#), among others, proved that SA would converge to a global optimum given the number of iterations converged to infinity. The pseudo-code of the method for a maximization problem is given in Algorithm 1.

At first, since it is a location-based method, an initial solution  $s_{start}$  is needed as well as the initial parameters  $\tau_0$  and  $L_0$  (temperature and number of iterations for  $\tau_0$ ). A solution  $s'$  that is in the neighborhood of  $s$  ( $\mathcal{N}_s$ ) is generated at random and the objective function values  $f$  of  $s$  and  $s'$  are compared. If the neighborhood solution  $s'$  is as good as or better than the current location  $s$  then the new solution becomes definitely the current location. Otherwise, the Metropolis criterion comes into play: if  $f(s')$  is worse than  $f(s)$ ,  $s'$  becomes the new location with a probability of  $e^{(f(s')-f(s))/\tau_k}$ . Consequently, small deteriorations and solutions found at a higher temperature  $\tau_k$  are more likely to be accepted. The probabilistic acceptance of inferior solutions is needed to prevent that the algorithm gets stuck in a local optimum. The temperature is held for a phase of  $L_k$  iterations; then, it is decreased according to a cooling schedule. Thus, the longer the search lasts the less likely the acceptance of a worse solution becomes. In this case, the cooling schedule is a step function. The procedure is repeated until a predetermined stopping criterion is fulfilled (e.g., reaching a certain number of iterations).

**Algorithm 1** Simulated Annealing (see Aarts and Eikelder 2002)

```

s_start, τ_0, L_0 ← initialize()
k ← 0
s ← s_start
repeat
  for l = 1, . . . , L_k do
    s' ← generate(N_s)
    if f(s') ≥ f(s) or e^( (f(s')-f(s))/τ_k ) > random[0, 1) then
      s ← s'
    end if
  end for
  k ← k + 1
  L_k ← setLength(L_{k-1})
  τ_k ← decreaseTemperature(τ_k)
until stopcriterion = true
    
```

<i>k</i>	temperature phase	<i>l</i>	neighborhood search iteration
<i>s, s'</i>	solutions	<i>L<sub>k</sub></i>	number of iterations for <i>k</i>
<i>τ<sub>k</sub></i>	temperature	<i>N<sub>s</sub></i>	neighborhood of <i>s</i>
<i>f</i>	objective function ( <i>here</i> : maximization problem)		

**Table 2** Negotiation as a prisoner’s dilemma

Agent 1	Agent 2	
	Greedy	Greedy
Greedy	C; C	A; D
Cooperative	D; A	B; B

$$A \geq B \geq C \geq D$$

$$2 * B \geq A + D$$

Because of the fact that SA not only accepts improvements, the acceptance ratio of the central optimizer is much higher than in a standard Hill Climbing heuristic (Klein et al. 2003; Fink 2006). Furthermore, in the first rounds of the procedure, the acceptance ratios are much higher than in last rounds as the annealing temperature leads to a decreased probability of a change of location.

In negotiations, the decisions are not made by a central optimizer that knows the objective function but rather by several decentralized entities that do not reveal their utility functions (see Assumption 4).

If an agent behaves in a SA manner, one could call it cooperative behavior. Imagine the following scenario: two agents negotiate with each other and an agent could achieve an improvement by a neighboring contract whereas the other agent would suffer a very marginal deterioration. Assuming an agent behavior similar to the SA behavior, the agent worse off could accept this contract and, thus, help the opposing agent.

Although this behavior is socially dominant, it is not individually dominant. Table 2 shows a game in which two agents can choose between cooperative and greedy (i.e., myopic) behavior; proportions have been deducted by means of computational experiments (see Klein et al. 2003; Lang and Fink 2012a).

This game is an instance of the famous prisoner's dilemma. The socially most beneficial outcome is undermined by individual optimizing behavior (Axelrod and Hamilton 1981; Tucker 1983). Although the social optimum is in the strategy set  $\{cooperative; cooperative\}$ , both agents have an incentive to deviate to greedy behavior. Also in the set  $\{greedy; cooperative\}$  (or vice versa), deviation from cooperative to greedy is individually rational. Finally, the resulting Nash equilibrium, which is the only one of this game, is  $\{greedy; greedy\}$  and simultaneously the social welfare minimum. Hence, a negotiation protocol needs appropriate mechanisms to enforce cooperative behavior and make it the individually reasonable strategy.

## 4.2 Basic Protocol

Simulated Annealing partly leads to acceptances of solutions that are worse than the currently best found. This can overcome local optima and, thus, facilitate finding global optima.

The protocol is based on the works of Klein et al. (2003, 2007) and Fink (2006). However, these works cannot comply with all of our requirements presented in Sect. 2.2.1 as discussed in the related work section (see Sect. 3). Having those drawbacks of previous works in mind, we propose a negotiation protocol which is supposed to achieve beneficial results according to the presented measures and to fulfill the presented requirements. The protocol itself is generic, i.e., it is not limited to a certain application. The main mechanism of the protocol is the application of acceptance quotas for a set of proposals per iteration. Although the protocol does not use probabilistic functions like the Metropolis criterion, it draws on this characteristic of the SA metaheuristic. That is why we call the protocol *Mediated Negotiation Protocol Based on Simulated Annealing* (MNP-SA). The pseudo code is given in Algorithm 2. The detailed implementation of the functions (sub-procedures) is presented in the next subsection.

At first, we need an initial draft contract  $c^a$ . We call this contract the active contract. In the following, the active contract draft is defined as the last overall accepted contract and current contract in force.

After the initialization, the negotiation procedure starts and is repeated until the  $T$ -th iteration. In every iteration  $t$ , we have to determine a quota  $p_t$  stating how many percent of the proposals have to be accepted. The acceptance quota declines over time ( $p_t \leftarrow p_1 * \beta^{t-1}$ , where  $\beta$  should be very close to 1). Subsequently,  $\rho$  mutations  $c^c$  of the active contract  $c^a$  are generated. Those mutations, along with the active contract, constitute the set of contract proposals. Afterwards, the agents decide whether they accept or reject the proposals, but are required to accept at least  $q_t = \lceil p_t * \rho \rceil$  contracts. If a contract candidate is accepted by a certain number of agents, it becomes a potential active contract for the next round and, thus, is added to the set of potential contracts *AccProposals*. If there is no such contract (*AccProposals* =  $\emptyset$ ), the active contract remains for the next iteration; if there is just one contract, this contract becomes the active contract in the next iteration; and if there are several contracts, the mediator selects one randomly. Thereafter, the process starts over and new proposals are generated using the (new) active contract.

**Algorithm 2** Mediated Negotiation Protocol Based on Simulated Annealing

---

```

 $c^a \leftarrow \text{GenerateInitialContract}()$  ▷ Generation of Initial Contract
for  $t = 1, \dots, T$  do
   $p_t \leftarrow p_1 * \beta^{t-1}$  ▷ Determination of Quotas
   $\text{AccProposals} \leftarrow \emptyset$ 
   $\text{Proposals} \leftarrow \{c^a\}$ 
  for  $l = \{1, 2, \dots, \rho\}$  do
     $c^c \leftarrow \text{ProposeMutation}(c^a)$  ▷ Proposal Generation
     $\text{Proposals} \leftarrow \text{Proposals} \cup \{c^c\}$ 
  end for
  for all  $j \in \mathcal{J}$  do
     $\vec{Z}_j \leftarrow \text{Accept}(\text{Proposals}, p_t)$  ▷ Agents' Decisions
  end for
  for all  $c^c \in \text{Proposals}$  do
    if  $\sum_{j \in \mathcal{J}} Z_j[c^c] \geq \text{AccThreshold}$  then ▷ Selection of Accepted Contracts
       $\text{AccProposals} \leftarrow \text{AccProposals} \cup \{c^c\}$ 
    end if
  end for
  if  $\text{AccProposals} \neq \emptyset$  then
     $c^a \leftarrow \text{RandomlySelect}(\text{AccProposals})$ 
  end if
end for
 $c \leftarrow c^a$ 

```

---

$t$	negotiation round	$p_t$	acceptance ratio
$c^a$	active contract	$\rho$	number of proposals
$c^c$	contract candidate	$Z_j$	agent response

---

Finally, after  $T$  iterations, the last active contract  $c^a$  becomes the final contract of the negotiation  $c$ .

### 4.3 Protocol Building Blocks

The protocol is subject to different configurations depending on the involved functions' actual implementation. The protocol is based on building blocks, i.e., different extensions are conceivable, which are presented subsequently and tested in the computational experiments in Sect. 7.

The first function of the negotiation protocol is the *generation of the initial contract*. Naturally, the mediator can only construct a random initial contract; we use this **mediator initialization** as preset procedure. Another approach might be a **prenegotiation**, in which the agents try to find a suitable starting contract that might be more just than a random guess. For instance, the mediator could propose a large number of random contracts and the agents could perform a runoff voting, i.e., they continue eliminating the candidates with the least votes until just one contract is left (possibly, the mediator may support this in the case of ties).

The *determination of the quotas* is given by the proposed formula. The formula ensures that—given a sufficiently high starting quota  $p_1$  is chosen—the quota scheme (the counterpart to the SA cooling schedule) passes through a lot of quota states. However, the protocol **without quotas** (i.e.,  $\beta = 0$ ) represents an alternative to **quotas**.

The next function of the protocol is the *proposal generation*. A proposal is generated by a single mutation, i.e., one item of the contract is altered. Again, the **mediator** could propose mutations randomly or the agents could propose mutations that they regard as desirable. Since the agents would propose individually improving contracts only and are tempted to accept just those, we decided to design the **agent proposal** version as a hybrid, i.e., the mediator randomly proposes some candidates and the agents propose the remaining ones successively (just one agent has its turn per round).

Although the agents basically can choose between accept and reject in the protocol (**binary logic**), they could also apply **tertiary logic** in the subprocedure of the *agents' decisions*. The agents respond  $Z_j[c^c] = 1$  for acceptance of a proposal and  $Z_j[c^c] = 0$  for rejection. In the tertiary case, they also could respond  $Z_j[c^c] = 0.5$ , if they are forced to an involuntary acceptance due to the quotas. If all agents state this case, the proposed contract would lead to a collective deterioration, which could be prevented. The tertiary logic leads to additional information revelation; however, the information is directed to the mediator and not necessarily to the other agents.

Finally, in the *selection of accepted contracts*, the protocol has to determine a certain acceptance threshold that is the minimal number of votes needed for being accepted overall. In the default case, we propose that an acceptance has to be made with **unanimity** ( $AccThreshold = J$ ). Another possible implementation is the use of **simple majority** ( $AccThreshold = \frac{J}{2}$ ).

## 5 A Mediated Negotiation Protocol Based on a Genetic Algorithm

Likewise as in the previous section, we develop and present the *Mediated Negotiation Protocol Based on a Genetic Algorithm* (MNP-GA). In this section, we describe the general Genetic Algorithm approach at first, then present the negotiation protocol and, finally, elucidate the protocol building blocks.

### 5.1 A Genetic Algorithm

A Genetic Algorithm (GA) simulates an evolutionary process that selects fitter solutions. Unsurprisingly, the first works to this topic originate in the computer simulation of biological evolution (Fogel 2006). In the seventies, the seminal work of Holland (1975) formed the term GA and laid the theoretical foundation for subsequent works (Michalewicz 1996). GAs are **population-based** metaheuristics, i.e., instead of considering one current solution, a GA pursues several solutions at the same time. A general GA pseudo-code is given in Algorithm 3.

Genetic algorithms constitute a population-based approach, they pursue several solutions simultaneously. Consequently, we need to initialize a first “breed”. This initial population is evaluated, i.e., fitness values are determined. Based on that, some of the fitter solutions are selected to become the parents of the next generation. Appropriate and common methods for selection are, e.g., the roulette-wheel selection or the tournament selection. In the roulette-wheel selection, the probability of a solution being selected as a parent is equal to its share of the sum of all fitness values of the population ( $f_{s'}/\sum_{s \in P} f(s)$ ) (Goldberg 1989). In the tournament selection, a subset

**Algorithm 3** A Genetic Algorithm after Michalewicz (1996)

---

```

 $t \leftarrow 0$ 
 $P(t) \leftarrow \text{initialize}()$ 
repeat
   $E(t) \leftarrow \text{evaluate}(P(0))$ 
   $t \leftarrow t + 1$ 
   $P(t) \leftarrow \text{select}(P(t - 1), E(t - 1))$ 
   $P(t) \leftarrow \text{recombine}(P(t))$ 
   $P(t) \leftarrow \text{mutate}(P(t))$ 
until stopcriterion = true

```

---

$P$	population	$E$	evaluation	$t$	iteration
-----	------------	-----	------------	-----	-----------

---

of the population is randomly chosen and ranked by its fitness: the higher a solution is ranked the higher its probability to be selected (Miller and Goldberg 1995). A selection building block is elitism, i.e., a few elitist solutions are selected which are directly added to the next generation without recombination or mutation (Goldberg 1989). The “chromosomes”, describing the solution (e.g., bit strings), are recombined between two parents. Finally, there could occur mutations that are random deviations from the parents’ genetic information.

## 5.2 Basic Protocol

A central aspect of GAs is that they consider several solutions simultaneously. In the course of a negotiation, there can be several contract proposals that have an identical number of acceptances. In the MNP-SA, just one of these contracts could be selected (e.g., randomly). On the other hand, one can follow a population-based approach by designing the protocol such that we do not have to decide between several candidates. Furthermore, evolutionary recombination is applied for proposal generation by analogy. The resulting MNP-GA is shown in Algorithm 4.

Likewise as in the common metaheuristic, we need an initial set of contract proposals. Again, the usage of quotas like presented in Sect. 4.2 is conceivable. That means, subject to their preferences and the determined quota, the agents have to accept a respective number of contracts from the proposal set. In contrast to the MNP-SA, the proposals are not selected from the neighborhood of an active contract. We have to determine an active contract draft  $c^a$  to ensure that high quality solutions are maintained and that there is a unique final result. The active contract is determined as in the MNP-SA, i.e., one of the contracts that fulfill the acceptance threshold is randomly chosen as active contract. Pretests have shown that, without such a contract, the solution quality is critically threatened to decline at the end of the negotiation. Afterwards, the elitists and the parents of the next population, that is the next set of contract proposals, are selected in accordance with the agents’ acceptance decisions. Elitism is an optional building block. Like mentioned above, the elitist contracts are directly adopted for the next population. The chromosomes of the parental contracts are recombined in some way, e.g., by taking the item values from one parent or the other at random (crossover recombination). The population size should be stable; so,



**Algorithm 4** Mediated Negotiation Protocol Based on a Genetic Algorithm

```

Proposals ← GenerateInitialPopulation()           ▷ Generation of Initial Population
for t = 1, ..., T do
    pt ← p0 * βt                                   ▷ Determination of Quotas
    for all j ∈ J do
        Zj ← Accept(Proposals, pt)
    end for
    ca ← selectActiveContract(Z1, ..., ZJ)
    Elitists ← selectElitists(Z1, ..., ZJ)           ▷ Elitists
    Parents ← select(Z1, ..., ZJ)                 ▷ Selection
    Children ← recombine(Parents)                   ▷ Recombination
    Children ← mutate(Children)                     ▷ Mutation
    Proposals ← ca ∪ Elitists ∪ Children
end for
c ← ca
    
```

$t$	negotiation round	$p_t$	acceptance ratio
$c^a$	active contract	$Z_j$	agent response

the number of children resulting from the recombination procedure depends on the number of elitists. The children contracts are mutated, i.e., their bits are flipped with a very small probability per bit (e.g., 1%). Finally, the active contract, the elitists, and the children contracts yield the proposals for the next round. This procedure is repeated until the  $T$ -th round. Then, the last active contract becomes the final contract.

5.3 Protocol Building Blocks

In the MNP-GA, there are also different building blocks which can be combined in some specific implementation.

To begin with, the *initial population* can again be generated randomly by the **mediator** (default) or determined by a **prenegotiation**. Also, the protocol can be applied **with or without quotas** (see Sect. 4.3). The usage of the *elitism* building block is also optional. With **elitism**, some contracts are selected as elitists and are directly adopted. A contract is selected as elitist if a certain number of acceptance decisions is reached (**unanimity** or **majority**).

Regarding the *selection*, we chose two different forms: **truncation selection** and **tournament selection**. We implemented the former as follows: firstly, we ordered the proposals by their number of acceptances and, secondly, selected a certain number of parents starting with the best ranked—those that are not selected are truncated. For the latter, we created several tournaments that consist of a randomly selected subset of the proposals. The subset is also ranked by the number of acceptances and then the winners of the tournaments are selected according to a probability distribution whereby the better ranked contracts have a higher probability to win.

Similarly to MNP-SA, the *recombination* can be conducted by the **mediator** or by **partly agent-based proposal generation** as well. Since there is a very large number of potential combinations, the mediator can heuristically generate contracts and the proposing agent chooses the best ones.

## 6 Discussion of Requirements

Besides the goal of achieving high quality contracts, we elucidated seven requirements in Sect. 2.2.1. In this section, we discuss to which extent the two protocols comply with these.

The first requirement refers to incentive compatibility (i.e., agents acting truthfully). We claim that the agents fulfill truthfulness—as proven by the following proposition:

**Proposition 1** *Given private information, the agents act truthfully.*

Let the proposal set be  $\{c_a, A, B, \bullet\}$  with  $c_a$  as current active contract. Considering some agent  $j$ , for proposal  $A$ , there are two relevant cases: (1)  $U_j(A) \geq U_j(c_a)$  and (2)  $U_j(A) < U_j(c_a)$ . In the first case,  $A$  is accepted anyways, as it is an improvement or at least equally favorable. For the second case, let us assume that  $c_a > A > B > c_\bullet \forall c_\bullet \in \{\bullet\}$  holds. If the agent just has to accept one proposal, he or she will choose  $c_a$ . However, if there are quotas to force the agent to accept a further contract, an agent would act truthfully, when he or she accepts the contract with the least deterioration compared to  $c_a$  which is  $A$ . Moreover, let us say the agent has a belief  $b_j(c)$  which determines the probability that a contract  $c$  is generally accepted by the other agents. Then, a risk-neutral agent would choose the least expected deterioration:  $\max\{b_j(A) * U_j(A), b_j(B) * U_j(B)\}$ . If  $b_j(A) > b_j(B)$ , choosing  $B$  over  $A$  could be the better choice, because the acceptance of the group is unlikelier. Thus, an unwanted contract can be chosen against the real preference, because the group is likely to reject it. However, we assumed private information (see Assumption 4). With private information, there is no knowledge about the other agents' preferences; thus, the beliefs are equal:  $b_j(A) = b_j(B) = \bar{b}_j$ . The decision function is now  $\max\{\bar{b}_j * U_j(A), \bar{b}_j * U_j(B)\}$  and leads to the truthful decision to prefer  $A$  over  $B$ , since homogeneous beliefs merely constitute a monotonic transformation.  $\square$

This result is transferable to the other agent action: the agent-based contract proposal.

The second requirement deals with individual rationality, i.e., an agent should never be worse off by participating in the negotiation. Both protocols, MNP-SA and MNP-GA, allow opting out of the negotiation at each point in the negotiation process. If an agent  $j$  has the opinion that he or she would obtain a negative utility in the end, he or she would opt out. Opting out means leaving the negotiation with a utility of zero which is equal to not participating ( $U_j = 0$ ). As a consequence, the utility gained in the negotiation is non-negative ( $U_j \geq 0$ ). In the next negotiation round, the remaining agents  $\mathcal{J} \setminus \{j\}$  can continue without this agent—the rules of the protocol stay unaffected. Starting over the negotiation is also an option, but not absolutely necessary.

The third requirement demands behavioral stability, i.e., the adopted strategy shall remain the same. In Proposition 1, we have shown that the best strategy is truthfulness. This claim holds also true if there are no quotas or different quota values. Also, the opting out of an agent does not affect the applied strategies. Consequently, the protocols can fulfill this requirement.

The fourth requirement implies that a protocol shall lead to a guaranteed agreement. The protocols make use of a contract draft, the active contract. If the negotiation is aborted, e.g., due to a communication breakdown, there is still the last overall agreed

contract. The protocol leads to success if there is at least one agent left that can gain a benefit from participating ( $\exists j | U_j > 0$ ). However, an issue disregarded so far is feasibility. In some applications, finding a feasible contract that is beneficial might be a non-trivial problem.

The fifth requirement concerns the simplicity of the optimal strategy, i.e., the optimal strategy should be easy to learn and to apply. Referring to Proposition 1, the optimal strategy is simple, since the agents just have to choose a certain number of proposals out of a set of proposals. Likewise, the proposal generation is simple. In the MNP-SA, we allowed the agents to mutate a single contract item; in the MNP-GA, the mediator provided a set of recombinations among which the agents can choose. Thus, the possible proposal set is congruent with the set of possible mutations (MNP-SA) or the set exogenously given by the mediator (MNP-GA). The agents just have to choose the best out of this superset. However, if the superset is not limited (e.g., unrestricted recombination), other rules are needed. For instance, there can be limited decision time for proposal generation. As a consequence, the agents would need smart procedures and their action would not be as obvious and as easy to learn anymore.

The sixth requirement is the demand for non-disclosure of private preference information (privacy). Besides considerations of the negotiation parties (see Sect. 2), Proposition 1 is based on the assumption of private information, which is why privacy should be maintained. In contrast to some other protocols in the literature, there is just information exchange between mediator and agents. The mediator might be supposed to be trustful, but what happens in the case of an untrustful mediator, i.e., if the agents' responses are unveiled? Beyond that, in some situations, there may be no non-biased mediator available (Lai and Sycara 2008; Lai et al. 2008). However, the proposed protocols do not really require an actual mediator agent and can be applied in a completely decentralized way as well. In a decentralized implementation, the agents could access a publicly available source code for the mediator software component and execute the mediator's tasks (generating and selecting contract proposals) in parallel by referring to a common random number service. In this case, relevant information is directly exchanged between the agents and the privacy issue would thus become even more relevant.

According to the design of the protocols, the responses by the agents only constitute specific ordinal relations but no cardinal utility values. Hence, it is not possible to deduce cardinal utility information of the agents by respective observations. However, the question remains whether it is conceivable to derive the underlying ordinal preference order from the unveiled information. In this regard we establish the following proposition for the two protocols, MNP-GA and MNP-SA:

**Proposition 2** *In the proposed protocols, a linear order cannot be deduced based on the decisions of the agents as long as the contract space  $\mathcal{C}$  is larger than  $T^2 \rho^2$ .*

In the course of the negotiation a mediator or agent can observe  $T * \rho$  proposals ( $T$ : number of rounds,  $\rho$ : number of proposals in each round). Let us suppose one could arrange those observed proposals, called set  $\mathcal{O}$ , perfectly, i.e., the proposals are ranked from 1 to  $T * \rho$ . As a binary relation on  $\mathcal{O}$  is a subset of the Cartesian product of  $\mathcal{O} \times \mathcal{O}$  (see Bouyssou and Vincke 2010), we would obtain a relation on  $\mathcal{O}$  with an upper bound for its cardinality of at most  $(T * \rho)^2 = T^2 \rho^2$ . The linear order on the

contract space  $\mathcal{C}$  (see Definition 3) has a lower bound for its cardinality of  $|\mathcal{C}|$ , which would be the equivalence (or indifference) relation. In this case, an agent is indifferent between every contract in  $\mathcal{C}$  such that  $|\mathcal{C}|$  observations could be sufficient for deriving a total order. Consequently, as long as  $|\mathcal{C}|$  is larger than  $T^2\rho^2$ , an agent’s preference relation cannot be deduced completely.  $\square$

For example, considering one million iterations ( $T$ ) and twenty proposals ( $\rho$ ) per iteration, the mediator could theoretically collect  $((10^6)^2 * 20^2 = 4 * 10^{14}$  concrete relations at most; however, the practically possible number would be drastically smaller. In our computational experiments (see Sect. 7), the decision space for an item is binary such that the contract space size for 100 items is  $2^I = 2^{100} > 10^{30}$ . A different example is the scheduling problem from Sect. 2.3, in which the contract space size for 100 items to be sorted is  $I! = 100! > 10^{157}$ . Consequently, the protocols just reveal a vanishingly small fraction of the contract space of typical complex negotiation applications—even if the upper bound of revelation is compared with the lower bound of specific contract relations.

Finally, the seventh requirement says that a protocol has to be scalable to many agents as well as many contract issues. The proposed protocols are not limited to a certain number of contract issues or agents and, therefore, also comply with this requirement. However, as we are going to show in Sect. 7, the more agents participate the larger the conflict of interests and the worse the achieved agreements may be. We expect that the same holds true for increasing numbers of items. Nevertheless, the computational results will show that the protocols are not only capable of comprising a large number of agents, but also can achieve beneficial results in terms of social welfare for them.

## 7 Computational Experiments

### 7.1 Simulation Scenario

In accordance to the problem model described in Sect. 2, we constructed a simulation testbed and conducted computational experiments. For the experiments, we generated 100 test instances for up to 10 agents ( $2 \leq J \leq 10$ ) and 100 contract issues ( $I = 100$ ). We supposed a binary decision space for all contract items ( $d_i \in \{0, 1\} \forall i \in \mathcal{I}$ ) which leads to a contract space size of  $2^I \approx 1.27 * 10^{30}$  (see Definition 3)—a number significantly too large to be explored exhaustively. Furthermore, we assumed the contract items to be pairwise interdependent and the utility function to be linearly

$$\text{additive: } U_j(c) = \sum_{p=1}^I \sum_{q=p}^I P_j(p, q) * d_p * d_q.$$

The experiments consider two cases for the generation of utility values ( $P_j$ ). Firstly, in the homogeneous case, all agents draw on the very same uniformly distributed utility value distribution:  $P_j(p, q) \sim \mathcal{U}(-100, 100)$ . Secondly, in the heterogeneous case, the agents are randomly assigned to one of four different agent types which draw on different beta distribution functions ( $\mathcal{B}$ ): (i) uniform:  $P_j(p, q) \sim (\mathcal{B}(1, 1) - 0.5) * 100$ , (ii) bell-shaped:  $P_j(p, q) \sim (\mathcal{B}(5, 5) - 0.5) * 100$ , (iii) left-skewed:  $P_j(p, q) \sim (\mathcal{B}(6, 2) - 0.75) * 100$ , and iv) right-skewed:  $P_j(p, q) \sim (\mathcal{B}(2, 6) - 0.25) * 100$ .

We implemented truthful agents in the simulation testbed. The rationale for this behavior has been discussed in Sect. 6. We tested the negotiation protocols using different policy building blocks configurations with one thousand, ten thousand, and hundred thousand iterations ( $T = \{1k, 10k, 100k\}$ ). In one iteration,  $\rho = 20$  proposals were generated. The computation for one problem instance took between half a second (1k rounds, two agents, simple building blocks) and half an hour (100k rounds, ten agents, several complex building blocks) using a 2.66 GHz processor (single threaded, no parallel processing of agents). We set the initial acceptance quotas rather high ( $p_1 = 0.75$  for  $J \leq 3$  and  $p_1 = 0.95$  for  $J > 3$ ) and used an annealing factor  $\beta = (p_T/p_1)^{1/(T-1)}$ . This factor results in a convex function that reaches a variety of different states. The quota starts with 0.75 or 0.95 and ends with a desired quota that we set to  $p_T \stackrel{!}{=} \frac{1}{\rho} = 0.05$ . Consequently, in the last round, just one contract has to be accepted ( $p_T * \rho = 0.05 * 20 = 1$ ).

Regarding the performance measurement, we approximated the Pareto frontier as well as the social welfare optimum by using a multi-objective evolutionary algorithm (MOEA, see, e.g., Zitzler and Thiele 1999). Precisely, we implemented a Memetic Algorithm, a hybrid version of a Genetic Algorithm, Hill Climbing, and SA to approximately compute the social welfare optimum  $c^*$ . Because of the complexity of the Pareto frontier computation, we computed it only for two and for three agents. We selected the best Nash product as well as the Kalai’s Egalitarian solution from the approximated set of Pareto-optimal solutions and computed the social welfare optima for two to ten agents separately.

Regarding the measurement of the performance, we used the Euclidean distance between a negotiation solution and the Pareto frontier / Nash solution / Kalai’s Egalitarian solution. Furthermore, we normalized the results of the different test instances by defining the claim points (the individually best solution which results from the application of the Memetic Algorithm) as 100%. Hence, the Pareto (P) / Nash (N) / Kalai’s Egalitarian (K) performance is computed as follows<sup>4</sup> (see Baarslag et al. 2013):

$$\{P(c), N(c), K(c)\} = \sqrt{\sum_{j \in \mathcal{J}} \left( \frac{U_j(c) - U_j(\{c^P, c^n, c^k\})}{U_j(c_j^{best})} \right)^2} \tag{1}$$

Concerning the social welfare (SW), we computed the measure as a ratio instead of an Euclidean distance:

$$SW(c) = \sum_{j \in \mathcal{J}} \frac{U_j(c)}{U_j(c^*)} \tag{2}$$

To test for statistical significance, we conducted Wilcoxon signed rank tests on the results. By doing so, we tested the hypothesis whether the better configuration is significantly superior (one-sided test). We set the threshold for a statistically significant result to a  $p$  value of 5%.

<sup>4</sup> See Sect. 2.2.2 for definitions of  $c^P$ ,  $c^n$ , and  $c^k$ ; furthermore,  $c_j^{best} = \arg \max_c U_j(c)$  (claim point) and  $c^P = \arg \min_{c^P \in \mathcal{P}} EuclideanDistance(c, c^P)$  (closest Pareto point).

**Table 3** Pareto, Nash, and Egalitarian results for the MNP-SA (homogeneous agents)

$J = 3$	Pareto (%)			Nash (%)			Kalai (%)		
	1k	10k	100k	1k	10k	100k	1k	10k	100k
No	49.66	52.07	51.31	59.63	61.42	61.16	58.43	60.51	60.20
Q	9.65	2.96	2.06	18.08	11.27	10.17	18.43	13.31	13.02
No+M	10.85	10.01	10.22	21.70	21.07	20.47	22.57	21.98	21.51
Q+M	11.71	10.54	10.52	21.77	21.76	20.35	22.54	23.00	21.59
No+A	47.96	52.35	49.71	57.92	61.95	59.33	57.21	61.09	58.71
Q+A	4.27	2.47	2.19	12.39	10.93	11.62	14.92	13.63	13.77
No+P	36.62	35.70	37.83	57.49	58.21	57.67	57.42	57.94	56.89
Q+P	9.69	3.03	2.32	18.88	11.44	10.05	19.54	14.23	12.41
Q+3	7.99	2.67	2.26	15.71	10.30	10.50	17.04	12.50	12.93
Q+P+A	4.62	2.45	2.24	13.80	11.08	10.56	16.03	13.87	13.34
Q+3+A	4.65	2.41	2.22	13.66	11.29	10.56	14.85	13.81	13.38
Q+P+3	6.95	2.85	2.23	14.24	11.04	10.78	15.40	13.37	13.90
Q+P+3+A	3.70	2.52	2.14	12.82	10.47	10.77	15.27	13.03	13.49

## 7.2 Results

### 7.2.1 Pareto, Nash, and Egalitarian Performance (Three Agents)

At first, we consider the case of homogeneous agents and take a closer look at the distance measures: the Pareto, Nash, and Egalitarian performance. We consider the outcomes close to the Pareto frontier as efficient, the outcomes close to the Nash product maximum as jointly efficient, and the outcomes close to the Egalitarian minimum maximization as socially fair.

The results (average values of the 100 problem instances) for the MNP-SA with three agents and 1,000, 10,000, as well as 100,000 iterations are given in Table 3. The first main result is that the configurations **with quotas (Q)** performed considerably better than **without quotas (No)** in terms of Pareto, Nash, and Egalitarian performance. For the case of no quotas, just the **majority rule (M)** can improve the results substantially, but it worsens the performance with quotas (Q+M). No+M is not statistically significantly better than Q+M in terms of Pareto outcomes ( $p$  values  $> 5\%$ ; see above). The majority rule is supposed to increase the joint acceptances of proposals, so do the quotas. The combination of both might lead to too many acceptances and, therefore, appears to overshoot the core of the quota idea.

The use of **agent-based proposals (A)** does not lead to significant Pareto improvement without quotas (No+A); however, with quotas (Q+A), they lead to significantly better Pareto results for 1k and 10k iterations (no significance for 100k). The Nash performance of Q+A is better compared to Q for 1k, equal for 10k, but worse for 100k. The Egalitarian measure is just better for 1k iterations (afterwards, it is statistically equivalent). The agent-based proposals shall lead to good moves, i.e., a smarter con-

tract proposal generation in each step. This seems to work out since it outperforms the basic version (Q) with few iterations (1k and 10k). However, due to individually optimizing actions, the proposals are presumably not the socially best possible. With more iterations, the need for good moves is smaller. Consequently, Q+A loses its advantage and the disadvantage seems to predominate. This becomes more evident for 1 million iterations (not shown in the table): Q has a mean Pareto distance of 1.90 %, whereas Q+A has a mean of 2.29 % (Nash: 9.05 & 10.64 %; Egalitarian: 12.39 & 13.85 %).

The findings of the **prenegotiation (P)** are vice versa: without quotas (No+P), there are Pareto, Nash as well as Egalitarian improvements; with quotas (Q+P), there is no statistically significant dominance for neither of the three measures. The intention of this configuration was that a more fair starting contract could lead to a more egalitarian final contract. However, the initial contract does not seem to have a considerable impact on the negotiation using this protocol for three agents.

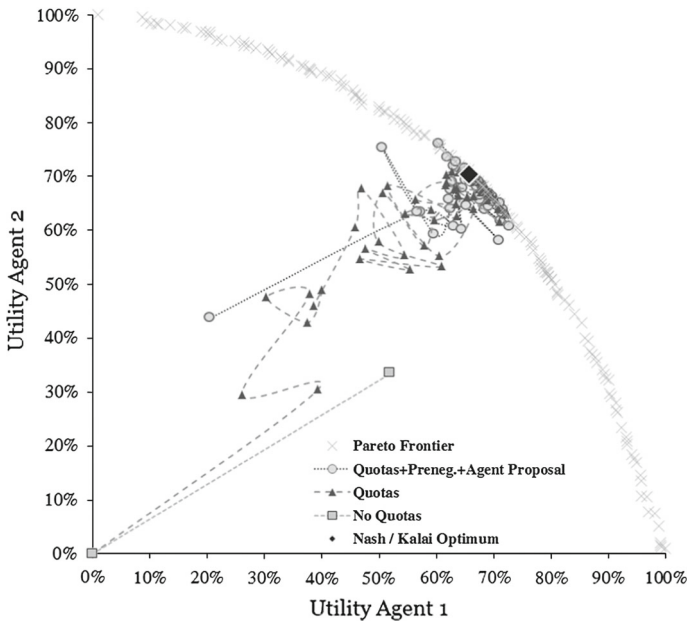
As mentioned before, the application of a **three-valued logic (3)** discloses more information and, therefore, should lead to better acceptance decisions by the mediator. On the other hand, the information revelation has to be considered as a disadvantage due to the requirement of privacy (see discussion in Sect. 6). The Pareto performance is slightly improved with 1k iterations, but this loses significance for more iterations.

Besides applying just one building block, the protocol can also make use of combinations of policies. The agent-proposal policy (Q+A) can be extended by additionally using three-valued logic (**Q+3+A**) or a prenegotiation (**Q+P+A**). Both extensions do not provide a significant increase in terms of Pareto or Egalitarian distance; however, the distance to the Nash solution is significantly smaller for 100k for both. Thus, applying an extra policy leads to a higher joint efficiency in the case of many iterations and might prevent some of the disadvantages discussed above. The combination of prenegotiation and three-valued logic (**Q+P+3**) does not result in any significant differences compared to Q+3 for the shown iteration numbers. However, with one million iterations, the outcomes of Q+P+3 are on average 1.51 % away from the Pareto frontier and significantly better than Q (1.90 %).

Finally, we can also apply all policies (**Q+P+3+A**), i.e., all policies except for the majority rule that appeared to be not favorable combined with quotas. Although the combination of two policies with quotas had little positive impact, the combination of three policies leads to substantially better Pareto results for the fewer numbers of iterations. However, the Nash and Egalitarian results are statistically equal to Q+A. With more iterations, the differences decrease and statistical significance fades. Though the three-valued logic policy did not provide a considerable gain on its own, it does in combination with other extensions.

Summarizing, the first main finding is that the configurations with quotas outperform the configurations without quotas significantly. Using the protocol's building blocks leads to better results for the fewer numbers of iterations. The second main finding is that the quota-based configurations—with 100,000 iterations—almost do not have any statistically significant differences compared to the protocol with quotas only. The only exceptions are the ones with majority rule (Q+M) (and the Nash outcome of Q+A). Nevertheless, on the one hand, some application scenarios can just allow few iterations, and, on the other hand, there are more complex problems than the





**Fig. 1** Exemplary visualization of the MNP-SA with two agents

ones used in the computational experiments. The consideration of a relatively small number of iterations is important if there are technical limitations: As the negotiation is distributed (e.g., between differently located companies), there may occur communication latencies. Furthermore, there may be a feasibility test which checks proposals for constraints. Depending on the complexity of the problem, such computations can take a substantial amount of time. Along with more complexity, finding near optimal outcomes becomes harder. Consequently, more iterations may be needed (e.g., a million or even more). If this is not possible due to runtime constraints, better moves are needed in one iteration. This can be provided by the proposed extensions of the basic protocol. Finally, the third main finding is that policies that are not very beneficial on its own, can be favorable in combination.

Figure 1 shows a characteristic visualization of negotiation histories of the MNP-SA (100,000 rounds) with two agents. Each dot represents the outcomes of the active contract after one thousand additional negotiation rounds (duplicates are left out). The black diamond in the middle depicts the Nash and Egalitarian optimum which are identical in this case.

The protocol configuration with quotas only (Q), represented by the triangular dots, approaches successively the Pareto frontier. The protocol allows deteriorations and, therefore, the negotiation history shows circular movement. After approx. 20,000 rounds, the protocol has reached the Pareto frontier, but still moves along it. Finally, an outcome ( $\{70\%; 66\%\}$ ) near to the Nash and Egalitarian optimum ( $\{66\%; 70\%\}$ ) is found in this specific instance. The configuration that includes quotas, a prenegotiation, and partly agent-based proposals (Q+P+A) is represented by the circular dots. Usually, the completely random starting contract may result in negative utilities for both agents

**Table 4** Pareto, Nash, and Egalitarian results for the MNP-GA (homogeneous agents)

$J = 3$	Pareto (%)			Nash (%)			Kalai (%)		
	1k	10k	100k	1k	10k	100k	1k	10k	100k
No	21.55	20.71	20.37	31.09	29.59	29.66	31.51	29.48	30.38
Q	19.34	21.37	21.89	29.24	31.45	30.79	29.60	31.93	31.25
No+M	20.31	22.07	19.27	29.41	32.45	29.51	30.37	33.20	29.80
Q+M	21.21	21.13	20.24	30.39	30.14	29.79	30.43	31.21	30.43
No+P	21.79	20.80	20.79	30.34	28.70	30.18	31.29	29.05	31.37
Q+P	20.19	20.53	21.48	29.02	29.62	31.33	29.76	30.11	31.68
No+T	22.38	21.08	22.27	32.03	29.96	32.60	32.49	30.49	33.40
Q+T	21.77	21.05	19.59	30.86	30.35	28.75	30.97	31.48	29.52
No+A	9.68	9.39	9.04	20.66	18.75	19.46	22.36	21.44	21.09
Q+A	9.57	9.23	9.08	19.18	20.64	20.00	21.18	22.54	21.75
No+E	35.23	27.28	16.60	43.67	37.22	26.02	43.29	37.00	26.04
Q+E	6.01	3.94	3.65	15.76	13.56	13.08	17.68	16.15	16.32
Q+P+T	21.40	21.23	21.02	30.62	29.49	29.56	31.10	30.14	30.10
Q+P+A	10.46	8.60	10.24	19.18	20.32	20.46	20.92	22.54	22.18
Q+E+P	6.05	4.10	3.28	14.25	13.16	12.70	16.53	15.61	16.34
Q+E+T	5.69	3.05	2.43	14.00	11.29	11.32	15.51	13.99	15.06
Q+E+A	3.13	2.30	2.38	15.06	13.07	12.73	18.12	16.21	16.33
Q+E+A+T	3.35	2.58	2.52	13.34	12.09	13.83	16.25	14.94	15.93
Q+E+P+A+T	3.20	2.93	2.42	14.05	12.94	10.78	17.03	16.43	14.08

so that both would opt out and get a utility of zero. However, the prenegotiation yields a better starting contract. The agent-based proposals lead to two patterns: firstly, the movement towards the frontier is more direct, but, secondly, the movement also has more spikes which favor just one of the agents—the spread is wider. The final contract ( $\{67\%; 69\%\}$ ) is very close to the Nash and Egalitarian optimum in this instance. Finally, represented by the square dots, the configuration without quotas (No) reaches a deadlock after very few rounds in which both agents reject to accept any further proposals; the final outcome is ( $\{52\%; 34\%\}$ ) and, thus, far away from efficiency.

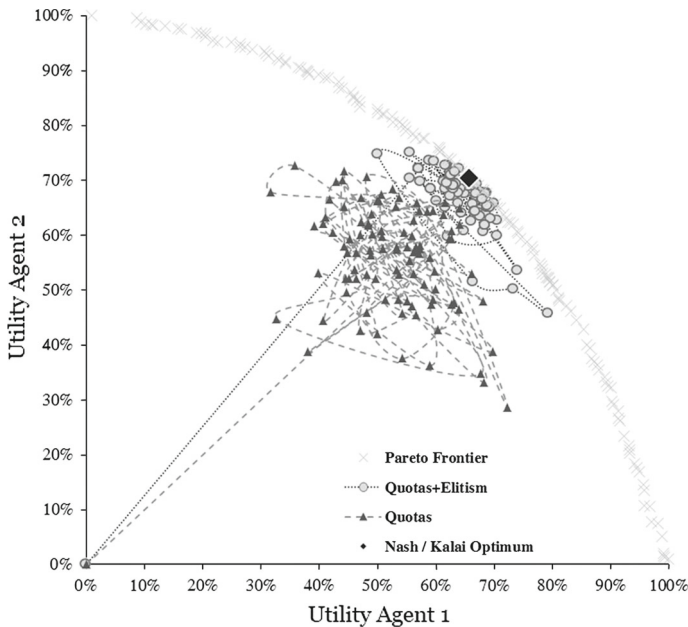
After having presented the results of the MNP-SA, we take a closer look at the population-based protocol, the MNP-GA. The computational results are shown in Table 4. First of all, in the population-based protocol, **quotas (Q)** on their own do not have a clear advantage to **no quotas (No)**. Both configurations perform about equally (Q is statistically slightly better for 1k iterations). The application of the **majority rule (M)** or the **prenegotiation (P)** does not affect the outcomes—the results of Q+M, No+M, Q+P, and No+P do not significantly differ from Q and No (exception: Egalitarian for No+M / No with 10k). When applying a **tournament selection (T)** instead of the simpler truncation selection, we obtain some significant differences, but the effects are still rather small and the scale stays about the same.

The Pareto, Nash, and Egalitarian results become substantially better when the proposals are partly **generated by the agents (A)**. However, there is still no statistical difference between quotas (Q+A) and no quotas (No+A). Differences between quotas and no quotas appear when the **elitism policy (E)** is applied. Without quotas (No+E), the results become drastically worse with few iterations, but slightly better with a lot of iterations. With quotas (Q+E), the results drastically improve with both, a few and a lot of iterations. This extension is the first one regarded so far that yields a considerable improvement. Consequently, we suppose that the population-based protocol has the most potential with quotas and elitism. Since the combinations **Q+P+T** and **Q+P+A** do not significantly differ from Q+T or Q+A, respectively, and are worse than Q+E, the experimental findings seem to support this hypothesis. Thus, we focus on configurations including quotas and elitism in the following.

Adding a prenegotiation (**Q+E+P**) does not significantly outperform the Q+E configuration of the protocol except for the Pareto performance with 100k iterations. Nevertheless, combining quotas and elitism with tournament selection (**Q+E+T**) leads to improvements. Except for the Pareto outcome with 1k iterations which fails statistical significance, Q+E+T can improve all three criteria for all three tested iteration numbers. Furthermore, the agent-based proposals (**Q+E+A**) achieve a further Pareto performance gain compared to Q+E+T; however, Q+E+A is most often outperformed by Q+E+T in terms of Nash and Egalitarian performance. Again, the agent-based proposal generation shows its advantage, good moves towards the Pareto frontier, but also its disadvantage, individually biased movements that may lead to unbalanced negotiation outcomes.

Finally, considering even higher degrees of configuration combinations, **Q+E+A+T** appears to be a compromise between Q+E+T and Q+E+A. Statistically, the Pareto performance is not distinguishable from Q+E+A and the Nash as well as Egalitarian performance is hardly distinguishable from Q+E+T, which significantly outperforms Q+E+A+T in the distance to the Nash solution with 100k iterations. Also, the combination of all but the majority rule, **Q+E+P+A+T**, represents a compromise between Pareto performance and Nash as well as Egalitarian performance. The Pareto performance is statistically equal to the results of Q+E+A+T; notably, in the case of 10k iterations, statistical significant dominance is just narrowly missed ( $p$  value = 5.1%). However, in terms of Nash and Egalitarian performance, the results tend to be better with Q+E+P+A+T for the cases with more iterations.

Summarizing, we firstly found that, all in all, the MNP-GA performs not as good as the MNP-SA in terms of distances to the Pareto frontier, Nash solution, and Egalitarian solution, but achieves just slightly worse outcomes. Nevertheless, like in the practice of metaheuristics, there may be problems that are better suited for the MNP-GA. The structure of the considered problem is rather well-behaved. If there are undefined regions in the problem space or severe feasibility constraints, a population-based protocol might have some advantages compared to a location-based approach. The second main finding is that the GA-based protocol suffers from much more instability. The protocol can hardly maintain good solutions and, therefore, final results are subject to a relatively large degree of randomness. This leads to the third main finding that only configurations that include quotas and elitism lead to satisfying results. The combi-



**Fig. 2** Exemplary visualization of the MNP-GA with two agents

nation of quotas and elitism can maintain good contracts better and, therefore, stays close to the Pareto frontier once it is reached.

Analogously to Fig. 1, typical negotiation histories of the MNP-GA with 100,000 rounds are shown in Fig. 2.

The configuration with quotas only, which performed similar to many other configurations, is represented by the triangular dots. The negotiation history shows the incapability of the protocol to maintain good solutions. Instead of moving towards the frontier more or less directly, the protocol jumps in the contract space. This protocol configuration fails to reach the Pareto frontier. The additional application of elitism (Q+E) leads to another picture. Now, the protocol approaches the frontier very fast and reaches it. However, as the protocol spends a lot of rounds close to the frontier, it takes long to reach it. More good solutions are maintained, but the outcomes are still relatively instable. The MNP-GA has much more duplicate active contracts (duplicates are left out in the figure) than the MNP-SA as well as more diversification. Finally, Q+E reaches a contract in this instance that is not directly at the Pareto frontier ( $\{68\%; 62\%\}$ ), but relatively close to the Nash and Egalitarian optimum ( $\{66\%; 70\%\}$ ). On the contrary, the configuration without elitism is far off ( $\{55\%; 60\%\}$ ) in this specific instance.

### 7.2.2 Heterogeneous Agents (Three Agents)

In this section, we analyze how the results are affected if the agents not only have different utility values but also draw on differently shaped utility distributions. As in

**Table 5** Comparison between homogeneous and heterogeneous agents (MNP-SA)

<i>J=3</i>	Pareto (%)			Nash (%)			Kalai (%)		
	<i>T=100k</i>	%	$\Delta$	Sig.	%	$\Delta$	Sig.	%	$\Delta$
No	42.56	-8.75	***	59.23	-1.93	°	54.57	-5.63	***
Q	2.28	0.22	**	17.84	7.67	†	20.69	7.67	†††
No+M	8.89	-1.33	**	26.88	6.41	†	26.87	5.36	†††
Q+M	9.37	-1.16	**	26.87	6.52	†	27.47	5.88	†††
No+A	42.06	-7.66	***	58.90	-0.43		54.30	-4.41	**
Q+A	2.55	0.36	°	18.82	7.20	‡	21.09	7.31	†††
No+P	31.97	-5.86	***	57.22	-0.45		54.50	-2.39	°
Q+P	2.55	0.23	*	17.56	7.51	‡	20.31	7.90	†††
Q+3	2.61	0.35	*	18.17	7.67	‡	21.06	8.13	†††
Q+P+A	2.65	0.40	°	19.56	9.00	††	21.52	8.18	†††
Q+3+A	2.77	0.56		18.88	8.32	††	21.29	7.91	†††
Q+P+3	2.37	0.13	**	18.13	7.35	†	20.59	6.69	†††
Q+P+3+A	2.51	0.37	*	18.83	8.06	†	21.88	8.39	†††

Heterogeneous case better at (significance level): \*\*\* 0.001, \*\* 0.01, \* 0.05, ° 0.1

Homogeneous case better at (significance level): ††† 0.001, †† 0.01, † 0.05, ‡ 0.1

the previous subsection, we tested three agents and determined the average distance to the Pareto frontier and Nash as well as Kalai's Egalitarian solution.

Table 5 shows the results for the MNP-SA with 100,000 negotiation rounds. For each measure, the table states the respective performance of the heterogeneous case (left column), the difference between the heterogeneous and homogeneous case (center column), and the significance obtained by a Wilcoxon signed rank test comparing the two cases (right column).

In terms of distance to the Pareto frontier, the outcomes for heterogeneous agents are better than for homogeneous agents for almost every protocol with a high significance. Although the well-performing configurations partly show a larger average distance than in the homogeneous case, the Wilcoxon signed rank test indicates that they are actually better except for Q+A, Q+P+A, and Q+3+A which do not achieve a 5% significance level. The larger mean can be explained with single outliers which dominate the mean if the configuration is otherwise close to the Pareto frontier. Despite the Pareto improvements, the final contracts are unfairer, i.e., the distribution of the outcomes are less uniform between the heterogeneous agents, as stated by the significance values for the distance to the Nash and Kalai solution. This could mean that some agents achieve very good results and, simultaneously, others do not reach as good outcomes. Some configurations show a better Nash and Kalai performance; however, we suppose that, since those configurations have underwent big Pareto improvements, they are naturally closer to the Nash or Kalai point without having fairer distributions of the negotiation success within the group.

The results for the heterogeneous agents experiments for the MNP-GA are shown in Table 6. The findings are almost the same as for the MNP-SA. The Pareto performance

**Table 6** Comparison between homogeneous and heterogeneous agents (MNP-GA)

<i>J=3</i>	Pareto (%)			Nash (%)			Kalai (%)			
	<i>T=100k</i>	%	Δ	Sig.	%	Δ	Sig.	%	Δ	Sig.
No	17.86	-2.52	**	33.57	3.91		32.25	1.87		
Q	18.16	-3.73	***	34.68	3.89		32.82	1.56		
No+M	17.24	-2.03	*	34.27	4.76		32.43	2.64	‡	
Q+M	18.15	-2.10	**	35.09	5.29	†	31.71	1.29		
No+P	17.03	-3.76	***	34.15	3.97		32.15	0.78		
Q+P	16.65	-4.84	***	32.07	0.74		30.72	-0.96		
No+T	18.39	-3.88	***	34.07	1.47		32.09	-1.31		
Q+T	17.20	-2.40	**	33.71	4.96		31.21	1.69		
No+A	8.17	-0.87	*	24.69	5.23	‡	25.55	4.46	††	
Q+A	7.81	-1.28	**	24.92	4.92		25.37	3.62	††	
No+E	11.08	-5.53	***	26.90	0.88		26.80	0.76		
Q+E	3.69	0.04	**	21.18	8.10	†††	22.30	5.97	†††	
Q+P+T	18.14	-2.88	**	34.49	4.93		32.51	2.41	†	
Q+P+A	8.94	-1.29	**	26.49	6.03	†	25.79	3.61	†	
Q+P+E	3.55	0.27	*	20.12	7.42	†	22.99	6.65	†††	
Q+E+T	2.72	0.28	**	19.19	7.87	††	21.90	6.85	†††	
Q+E+A	2.49	0.11	***	21.03	8.31	†††	24.11	7.78	†††	
Q+E+A+T	2.69	0.18	***	20.44	6.61	†	23.08	7.16	†††	
Q+E+P+A+T	2.68	0.26	**	20.53	9.75	†††	23.48	9.40	†††	

Heterogeneous case better at (significance level): \*\*\* 0.001, \*\* 0.01, \* 0.05, ° 0.1  
 Homogeneous case better at (significance level) ††† 0.001, †† 0.01, † 0.05, ‡ 0.1

is better throughout all configurations. The Nash and Kalai performance is worse for configurations that achieve a good Pareto outcome which supports the above-mentioned presumption. For most MNP-GA configurations, there is no significant difference between the two cases with regard to the Nash and Kalai results.

Concluding, in our experiments, the agents tend to be better off in terms of Pareto performance in the heterogeneous case, but tendentially worse off in terms of Nash and Kalai performance. This means that more heterogeneous agents are able to find better and more efficient contracts, but the outcomes are more unbalanced than in the case of rather homogeneous agents.

### 7.2.3 Social Welfare (Up to Ten Agents)

As mentioned before, the Pareto frontier, Nash solution, and Egalitarian solution are hard to compute for many agents. Therefore, we draw on the social welfare measure to quantify the performance of the protocols for many agents. We approximated the social welfare optimum for the 100 problem instances using a Memetic Algorithm which combines an Evolutionary Algorithm with Hill Climbing and Simulated Annealing.




Fig. 3 Social welfare: MNP-SA with 2–10 agents

In the following, 100 % represents the approximated social welfare optimum. In this section, the utility values are supposed to be  $P_j(p, q) \sim \mathcal{U}(-100, 100)$  again.

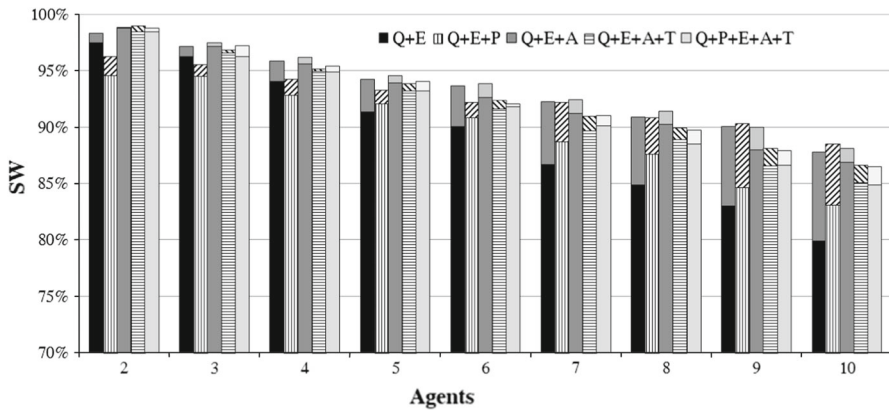
For the MNP-SA, the average social welfare outcomes for five configurations are shown in Fig. 3. The figure comprises two to ten agents. The lower bars show the results for 10,000 iterations, whereas the upper bars show the improved results for 100,000 iterations. The legend refers to the shades and patterns of the lower bars.

As the results regarding Pareto efficiency have already shown, the protocol configurations deliver very good solutions for small numbers of agents (e.g.,  $J = 2$ ;  $T = 100k$ : 99.0–99.2 %). With few agents, the differences between 10k and 100k iterations as well as between the configurations are relatively small. Certainly, with more agents, the gaps between the different numbers of iterations become significantly larger [e.g., Q assuming  $J = 10$ : 72.8 % (10k) & 78.3 % (100k)]. The overall decline with more agents is not surprising, as more agents result in stronger conflicts of interest among them (see Price of Anarchy: Roughgarden and Tardos 2007). When there are plenty of agents, few iterations result in a disparity of configuration performance. Applying more iterations counteracts the disparities of the configurations and leads to similar results. Furthermore, applying more iterations slows down the welfare decline caused by the arising conflicts between the agents. As shown, more elaborate protocol configurations tend to perform better. Applying a prenegotiation, three-valued logic, or agent-based proposing usually yield significantly improved outcomes for a large number of agents and 10k iterations. However, this trend loses some significance given there are 100k negotiation rounds. In the analysis of Pareto, Nash, and Egalitarian outcomes, there was no clear benefit from using many building blocks within the MNP-SA; however, here, the results show that combining these building blocks is favorable, as they generally require less iterations in more conflicting scenarios.

The MNP-GA has shown, with regard to the Pareto, Nash, and Egalitarian measures, for three agents slightly weaker outcomes. Figure 4 shows the social welfare performance of five configurations of the MNP-GA for two to ten agents.

Again, the MNP-GA does not achieve results as good as the outcomes of the MNP-SA for few agents (e.g.,  $J = 2$ ;  $T = 100,000$ : 96.2–99.0 %). Nevertheless,





**Fig. 4** Social welfare: MNP-GA with 2–10 agents

the MNP-GA is less affected by the increasing conflict than the MNP-SA—the performance for a large number of agents is much better. Since the MNP-GA is population-based, it can explore several regions of the contract space simultaneously. More agents can exacerbate the problem drastically (more complex contract space). A location-based approach might become stuck in an area, whereas a population-based approach can more easily diversify the search in other regions by recombination. Notably, the disparity of the results depending on different configurations as well as on different numbers of iterations decreases. Regarding single building blocks, tournament selection does not yield a statistically significant improvement assuming a large number of agents, although it has been shown beneficial for few agents in the Pareto analysis. With few iterations (10k), conducting a prenegotiation results in a significantly smaller improvement than agent-based proposals. With more iterations (100k), as there is no statistical difference, both configurations seem to perform similarly.

## 8 Conclusions and Future Work

Operational planning frequently comprises challenging, complex problems. Furthermore, corporate planning faces increasingly intercompany issues. For instance, companies are more and more connected within supply chains. Because of decentralized decision making, planning not only has to deal with complex problems, but also with strategies and interactions of self-interested parties. At this, negotiations are a powerful mechanism which is supposed to reconcile conflicting interests.

Since a lot of planning problems are frequently reoccurring, there is a large potential for automation of decision making by suitable algorithms. Such algorithms are, e.g., metaheuristics which are supposed to find near-optimal solutions for centralized problems in a given—preferably short—time. In this study, we present two negotiation protocols for automated group decision making by means of automated negotiation between software agents. The protocols are inspired by certain properties and procedures of metaheuristics, namely Simulated Annealing and Genetic Algorithms. Each of the two protocols consists of several policy building blocks which can be optionally applied.

For evaluation purposes, we conducted extensive computational experiments using a simulation testbed. The data shows that both protocols—given an eligible configuration—are able to reach high quality and fair outcomes. For a small number of agents, the protocols achieve nearly optimal social welfare results. With a rising number of agents, the outcomes become worse (using the same fixed number of iterations); nevertheless, both protocols still result in a beneficial performance considering the fact that more agents lead to a more severe conflict of interests. Generally, the findings provide insights regarding an effective configuration of the building blocks (e.g., quotas) and a reasonable parameterization (e.g., concerning the number of negotiation rounds).

Furthermore, this study has elaborated and emphasized seven requirements for protocol design. In general, the proposed protocols satisfy these demands adequately (with some limitations under particular circumstances).

The validity of the results is limited by the assumptions of the computational experiments. Generally, computational studies suffer from a lacking ability of generalization. Regarding different scenarios, other contract spaces may lead to different results. Especially in the case of agent preferences which partly depend on each other, e.g., in seller-buyer negotiations, the findings are hardly transferable. In this research, we have focused on exploiting win-win opportunities and ignored scenarios with zero-sum game characteristics. Moreover, we neglected feasibility issues of contracts which exacerbate proposal generation.

Future work includes the analysis of the performance of the protocols for complex contract spaces beyond the scenario regarded here. We will study certain negotiation problems including real-world applications such as intercompany machine scheduling or energy management. At this, we are going to take a closer look into side constraints which limit the feasible contract space and make the involved decisions more difficult. Furthermore, the impact of the number of contract items as well as the number of iterations is going to be analyzed in more detail. Future work may also incorporate enhancements of the protocols' design such as further policy building blocks. Finally, we are going to work out an architecture for related negotiation systems. The purpose of such a system is, firstly, to facilitate the negotiation parties in finding sophisticated configurations and parameterizations—since the choice of protocol can be a negotiation problem in its own right—and, secondly, support the agents in the execution of the negotiation.

## References

- Aarts E, Eikelder HT (2002) Simulated annealing. In: Pardalos P, Resende MGC (eds) Handbook of applied optimization. Oxford University Press, New York, pp 209–221
- Axelrod R, Hamilton W (1981) The evolution of cooperation. *Science* 211(4489):1390–1396
- Baarslag T, Fujita K, Gerding E, Hindriks K, Ito T, Jennings NR, Jonker CM, Kraus S, Lin R, Robu V, Williams CR (2013) Evaluating practical negotiating agents: results and analysis of the 2011 international competition. *Artif Intell* 198:73–103
- Bergstrom TC, Varian HR (1985) When do market games have transferable utility? *J Econ Theory* 35(2):222–233
- Bichler M, Kersten G, Strecker S (2003) Towards a structured design of electronic negotiations. *Group Decis Negot* 12(4):311–335
- Bichler M, Gupta A, Ketter W (2010) Designing smart markets. *Inf Syst Res* 21(4):688–699

- Binmore K, Vulkan N (1999) Applying game theory to automated negotiation. *Netnomics* 1(1):1–9
- Blum C, Roli A (2003) Metaheuristics in combinatorial optimization: overview and conceptual comparison. *ACM Comput Surv* 35(3):268–308
- Boutilier C, Shoham Y, Wellman MP (1997) Economic principles of multi-agent systems. *Artif Intell* 94(1–2):1–6
- Bouyssou D, Vincke P (2010) Binary relations and preference modeling. In: Bouyssou D, Dubois D, Pirlot M, Prade H (eds) *Decis Mak Process Concepts Methods*. Wiley-ISTE, London, pp 49–84
- Cerný V (1985) Thermodynamical approach to the traveling salesman problem: an efficient simulation algorithm. *J Optim Theory Appl* 45(1):41–51
- Collins J, Gini M, Mobasher B (2002) Multi-agent negotiation using combinatorial auctions with precedence constraints. Technical Report 02–009, University of Minnesota, Minneapolis, MN, USA
- Conitzer V (2010) Making decisions based on the preferences of multiple agents. *Commun ACM* 53(3): 84–94
- Conitzer V, Sandholm T (2004) Self-interested automated mechanism design and implications for optimal combinatorial auctions. In: Breesee JS, Feigenbaum J, Seltzer MI (eds) *Proceedings of the 5th ACM conference on electronic commerce (ACM-EC 2004)*, pp 132–141
- Fink A (2006) Supply chain coordination by means of automated negotiations between autonomous agents. In: Chaib-draa B, Müller J (eds) *Multiagent based supply chain management (Studies in Computational Intelligence, Vol. 28)*. Springer, Berlin, pp 351–372
- Fogel D (2006) *Evolutionary computation: toward a new philosophy of machine intelligence*, 3rd edn. Wiley-IEEE Press, Piscataway
- French S (1986) *Decision theory: an introduction to the mathematics of rationality*. Horwood, Chichester
- Fujita K, Ito T, Klein M (2010) Representative based multi-round protocol based on revealed private information for multi-issue negotiations. *Multiagent Grid Syst* 6(5–6):459–476
- Fujita K, Ito T, Klein M (2010) Secure and efficient protocols for multiple interdependent issues negotiation. *J Intell Fuzzy Syst* 21(3):175–185
- Fujita K, Ito T, Klein M (2014) Efficient issue-grouping approach for multiple interdependent issues negotiation between exaggerator agents. *Decis Support Syst* 60:10–17
- Goldberg DE (1989) *Genetic algorithms in search, optimization, and machine learning*. Addison-Wesley Professional, Boston
- Granville V, Krivanek M, Rasson JP (1994) Simulated annealing: a proof of convergence. *IEEE Trans Pattern Anal Mach Intell* 16(6):652–656
- Harsanyi J (1955) Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *J Polit Econ* 63(4):309–321
- Hattori H, Klein M, Ito T (2007) Using iterative narrowing to enable multi-party negotiations with multiple interdependent issues. In: Durfee EH, Yokoo M, Huhns MN, Shehory O (eds) *Proceedings of the sixth international joint conference on autonomous agents and multiagent systems (AAMAS 2007)*
- Holland JH (1975) *Adaptation in natural and artificial systems*. University of Michigan Press, Ann Arbor
- Hombberger J (2010) Decentralized multi-level uncapacitated lot-sizing by automated negotiation. *4OR* 8(2):155–180
- Hombberger J (2011) A generic coordination mechanism for lot-sizing in supply chains. *Electron Commer Res* 11(2):123–149
- Hombberger J (2012) A  $(\mu, \lambda)$ -coordination mechanism for agent-based multi-project scheduling. *OR Spectr* 34(1):107–132
- Hurwicz L (1945) The theory of economic behavior. *Am Econ Rev* 35(5):909–925
- Hurwicz L (1973) The design of mechanisms for resource allocation. *Am Econ Rev* 63(2)
- Imai H (1983) Individual monotonicity and lexicographic maxmin solution. *Econometrica* 51(2):389–401
- Ito T, Klein M, Hattori H (2008) A multi-issue negotiation protocol among agents with nonlinear utility functions. *Multiagent Grid Syst* 4(1):67–83
- Jennings NR, Sycara K, Wooldridge M (1998) A roadmap of agent research and development. *Auton Agents Multi-Agent Syst* 1(1):7–38
- Jennings NR, Faratin P, Lomuscio AR, Parsons S, Sierra C, Wooldridge M (2001) Automated negotiation: prospects, methods and challenges. *Group Decis Negot* 10(2):199–215
- Kalai E (1977) Proportional solutions to bargaining situations: interpersonal utility comparisons. *Econometrica* 45(7):1623–1630
- Kalai E, Smorodinsky M (1975) Other solutions to Nash's bargaining problem. *Econometrica* 43(3):513–518

- Kallrath J (2002) Planning and scheduling in the process industry. *OR Spectr* 24(3):219–250
- Kaneko M (1976) Note on transferable utility. *Int J Game Theory* 5(4):183–185
- Kirkpatrick S, Gelatt CD, Vecchi MP (1983) Optimization by simulated annealing. *Science* 220:671–680
- Klein M, Faratin P, Sayama H, Bar-Yam Y (2003) Negotiating complex contracts. *Group Decis Negot* 12(2):111–125
- Klein M, Faratin P, Sayama H, Bar-Yam Y (2007) Negotiating complex contracts. MIT Sloan School of Management Working Paper No. 4196–01, Massachusetts Institute of Technology, Cambridge, <http://ssrn.com/paper=290147>
- Kraus S (1996) An overview of incentive contracting. *Artif Intell* 83(2):297–346
- Kraus S (1997) Negotiation and cooperation in multi-agent environments. *Artif Intell* 94(1–2):79–97
- Kraus S (2001) Automated negotiation and decision making in multiagent environments. In: Luck M, Marik V, Stepankova O, Trappl R (eds) *Multi-agents Syst Appl*, vol 104. Springer, New York, pp 150–172
- Kraus S, Wilkenfeld J, Zlotkin G (1995) Multiagent negotiation under time constraints. *Artif Intell* 75(2):297–345
- van Laarhoven P, Aarts E (1987) *Simulated annealing: theory and applications*. D. Reidel Publishing, Dordrecht
- Lai G, Li C, Sycara K, Giampapa J (2004) Literature review on multi-attribute negotiations. Technical Report CMU-RI-TR-04-66, Carnegie Mellon University, Pittsburgh, PA, USA
- Lai G, Li C, Sycara K (2006) Efficient multi-attribute negotiation with incomplete information. *Group Decis Negot* 15(5):511–528
- Lai G, Sycara K (2008) A generic framework for automated multi-attribute negotiation. *Group Decis Negot* 18(2):169–187
- Lai G, Sycara K, Li C (2008) A decentralized model for automated multi-attribute negotiations with incomplete information and general utility functions. *Multiagent Grid Syst* 4(1):45–65
- Lang F, Fink A (2012) A quota-based multi-agent negotiation protocol for complex contracts. In: Filipe J, Fred ALN (eds) *ICAART 2012—Proceedings of the 4th international conference on agents and artificial intelligence*, volume 2—agents, SciTePress, pp 113–118
- Lang F, Fink A (2012) Collaborative job processing on a single machine: a multi-agent weighted tardiness problem. In: van der Hoek W, Padgham L, Conitzer V, Winikoff M (eds) *Proceedings of the 11th international conference on autonomous agents and multiagent systems (AAMAS 2011)*, vol 3, pp 1417–1418
- Lang F, Fink A (2012) Collaborative single and parallel machine scheduling by autonomous agents. In: Smari WW, Fox GC (eds) *2012 International conference on collaboration technologies and systems (CTS 2012)*, IEEE, pp 76–83
- Lang F, Fink A (2012) Negotiating in dynamic environments: time-efficient automated negotiations by means of combinatorial auctions. *Evol Syst* 3(3):189–201
- Lang F, Fink A (2013) Information sharing in interorganizational systems—a story about falsehood, greed, and privacy. In: Martinovski, B (ed) *Proceedings of the 13th meeting on Group Decision and Negotiation (GDN 2013)*, pp 319–323
- Lomuscio A, Wooldridge M, Jennings NR (2003) A classification scheme for negotiation in electronic commerce. *Group Decis Negot* 12(1):31–56
- Metropolis N, Rosenbluth A, Rosenbluth M, Teller A, Teller E (1953) Equation of state calculations by fast computing machines. *J Chem Phys* 21(6):1087–1092
- Michalewicz Z (1996) *Genetic algorithms + data structures = evolution programs*, 3rd edn. Springer, Berlin
- Miller BL, Goldberg DE (1995) Genetic algorithms, tournament selection, and the effects of noise. *Complex Syst* 9(3):193–212
- Nash JF (1950) The bargaining problem. *Econometrica* 18(2):155–162
- Neumann D (2007) *Market engineering: a structured design process for electronic markets*. Universitätsverlag Karlsruhe, Karlsruhe
- von Neumann J, Morgenstern O (1944) *Theory of games and economic behavior*. Princeton University Press, Princeton
- Ockenfels A, Roth A (2002) The timing of bids in internet auctions: market design, bidder behavior, and artificial agents. *AI Mag* 23(3):79–87
- Raiffa H (1982) *The art and science negotiation*. Harvard University Press, Cambridge
- Raiffa H, Richardson J, Metcalfe D (2002) *Negotiation analysis: the science and art of collaborative decision making*. Harvard University Press, Cambridge

- Rief D, Dintner C (2010) Negotiation for cooperation in logistics networks: an experimental study. *Group Decis Negot* 19(3):211–226
- Rosenschein JS, Zlotkin G (1994) Designing conventions for automated negotiation. *AI Mag* 15(3):29–46
- Roth A (1979) An impossibility result concerning n-person bargaining games. *Int J Game Theory* 8(3):129–132
- Rothkopf M, Teisberg T, Kahn E (1990) Why are Vickrey auctions rare? *J Polit Econ* 98(1):94–109
- Roughgarden T, Tardos É (2007) Introduction to the inefficiency of equilibria. In: Nissan N, Roughgarden T, Tardos É (eds) *Algorithmic game theory*. Cambridge University Press, Cambridge, pp 443–458
- Sandholm T (1999) Distributed rational decision making. In: Weiß G (ed) *Multiagent systems: a modern introduction to distributed artificial intelligence*. MIT Press, Cambridge, pp 201–258
- Satterthwaite MA (1975) Strategy-proofness and arrow's conditions: existence and correspondence theorems for voting procedures and social welfare functions. *J Econ Theory* 10(2):187–217
- Shoham Y, Leyton-Brown K (2009) *Multiagent systems: algorithmic, game-theoretic, and logical foundations*. Cambridge University Press, Cambridge
- Ströbel M, Weinhardt C (2003) The Montreal taxonomy for electronic negotiations. *Group Decis Negot* 12(2):143–164
- Sycara K (1998) Multiagent systems. *Artif Intell* 19(2):79–92
- Tucker AW (1983) The mathematics of Tucker: a sampler. *Two-Year Coll Math J* 14(3):228–232
- Tung HW, Lin RJ (2005) Automated contract negotiation using a mediation service. In: Müller G, Lin KJ (eds) *Seventh IEEE international conference on E-commerce technology (CEC 2005)*, IEEE, pp 374–377
- Varian HR (2010) *Intermediate microeconomics: a modern approach*, 8th edn. W. W. Norton & Company, New York
- Vetschera R (2013) Negotiation processes: an integrated perspective. *EURO J Decis Process* 1(1–2):135–164
- Vickrey W (1961) Counterspeculation, auctions, and competitive sealed tenders. *J Financ* 16(1):8–37
- Vulkan N (1999) Economic implications of agent technology and e-commerce. *Econ J* 109(453):67–90
- Zitzler E, Thiele L (1999) Multiobjective evolutionary algorithms: a comparative case study and the strength Pareto approach. *IEEE Trans Evol Comput* 3(4):257–271
- Zlotkin G, Rosenschein JS (1996) Mechanism design for automated negotiation, and its application to task oriented domains. *Artif Intell* 86(2):195–244