

A Survey on One-To-Many Negotiation: A Taxonomy of Interdependency

Tamara C.P. Florijn^{1,2}, Pinar Yolum² and Tim Baarslag^{1,3}

¹Centrum Wiskunde & Informatica (CWI)

²Utrecht University

³Eindhoven University of Technology

{tamara.florijn, t.baarslag}@cwi.nl, p.yolum@uu.nl

Abstract

One-to-many negotiations are widely applied in various domains, contributing to efficient resource allocation and effective decision making. This wide variety of applications also brings a wide variety of implemented protocols, terminology and utility functions, which makes it hard to compare and improve strategies using existing solutions. We introduce a meta-model of negotiations, which characterizes almost all one-to-many negotiation research, bringing a unified description of the negotiations. This meta-model allows us to identify different classes of interdependency based on utility functions. We show how existing one-to-many negotiations are related to each other, finding new insights and identifying knowledge gaps. We suggest that a general utility function framework and benchmark scenarios for one-to-many negotiations could accommodate future advancement in this field.

1 Introduction

Negotiation is a complex and time-consuming process to reach agreements between multiple parties, with applications ranging from energy grid networks [Brazier *et al.*, 2002] to supply chain management [Wong and Fang, 2010; Ruofei, 2016; Mohammad *et al.*, 2019]. The research area of *automated negotiation* explores how to make this process more efficient and reach better agreements by employing computerized strategies in the form of *agents* that negotiate with each other.

One particularly promising application domain of negotiation agents is e-commerce, in which one agent negotiates with multiple others (so called *one-to-many negotiations*), in an effort to buy and sell products at the best price. For example, consider a buyer (the ‘one’) that aims to buy strawberries, bananas and pears from multiple fruit sellers (the ‘many’). The negotiations with the fruit sellers can influence each other greatly; for example, to prepare a smoothie, the buyer might require equal amounts of strawberries, bananas and pears, and thus all negotiations should be coordinated to make sure that the (expected) total amount of each fruit remains equal over all deals.

One-to-many negotiations often involve such interdependencies between the different negotiations, in which the goals in one negotiation are dependent on the actions in other ones, restricting the space of favorable offers. The coordination of these different negotiations is hard because (1) the outcome depends on uncertain factors such as the future actions of the opponents, and (2) the agent faces an overwhelming number of options to choose from, as each negotiation can present numerous possible actions that can be combined with those from other negotiations.

To deal with these interdependencies, varied types of research have been published in one-to-many negotiation. One of the first papers in the field appears in Rahwan *et al.* [2002], introducing a protocol where agents are assigned to strike a single deal with one of multiple opponents. Much of the following work expands on this with new agent strategies [Nguyen and Jennings, 2004; Williams *et al.*, 2012; Mansour and Kowalczyk, 2015], for example with the ability to cancel a prior deal [Nguyen and Jennings, 2005; An *et al.*, 2008], or requiring a confirmation before the deal becomes binding [An *et al.*, 2009]. This research trend has continued to address increasingly difficult settings, mirroring real world applications such as procurement and supply-chain management [Mohammad, 2021; Baarslag *et al.*, 2021; Najjar *et al.*, 2021; Li *et al.*, 2024].

A key difference between these one-to-many negotiations is often how research regards the relationship between the ‘one’ and the ‘many’. Their interdependency can come in different forms, including constraints on how many deals may be formed [Bagga *et al.*, 2021; Alrayes *et al.*, 2018] or how much money is spent in total [Mansour and Kowalczyk, 2015; Najjar *et al.*, 2021].

To shed some light on the different types of relationships within the one-to-many negotiation literature, this paper provides an integrative survey on one-to-many negotiation and their interdependencies. Our contributions consist of (1) a method to quantitatively analyze one-to-many negotiations based on the utility function (see Figure 2), and (2) a taxonomy of how existing one-to-many negotiations are related, classifying the research by their interdependencies.

To do so, we survey all recent papers on one-to-many negotiation through the lens of an *agent designer*. That is, we focus on the challenges that agent designers face, guided by their evaluation method. We start by defining a meta-model

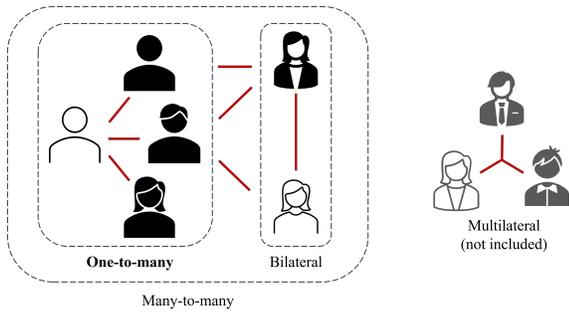


Figure 1: A visual representation of the difference between one-to-many, many-to-many negotiation and multilateral negotiation.

(Section 4) that allows us to link up existing one-to-many negotiation research. Following that, we extract relevant information from the meta-model that forms the input of the utility function. This allows us to design a new taxonomy based on the properties of the utility function to classify the studies in Section 5. Researchers in the field can use this taxonomy to recognize similar problems and exploit common properties within the class of a negotiation setting, and identify closely related papers. Furthermore, it can guide the research community in identifying unexplored research topics, which we elaborate on in the final sections of the paper.

2 One-To-Many Negotiations

In one-to-many negotiations, one *center* agent α negotiates with m opponents $\beta_1, \beta_2, \dots, \beta_m$ in multiple bilateral *subnegotiations*, for example in fruit procurement. A negotiation can be multi-issue (e.g. the number, price and quality of fruits) or single issue (e.g. the number of one type of fruit). This choice depends on what issues influence the *utility* of an agreement, typically expressed as a function u mapping all possible outcomes to a real number ranging from 0 to 1.

The agreement that the center agent achieves does not always consist of only one deal (*single deal*) but can also consist of multiple deals that it struck with multiple opponents (*multi-deal*), e.g. a fruit buyer that collects fruit from multiple sellers. The set of all possible outcomes (that is all agreements including non-agreement \emptyset) is called the outcome space Ω . If the outcome space is not shared, Ω consists of subdomains $\Omega_1, \Omega_2, \dots, \Omega_m$ for each opponent $\beta_1, \beta_2, \dots, \beta_m$.

In the context of multi-deal negotiation, ambiguity can arise from the term “agreement”. Specifically, it may refer to either a bilateral agreement between the center agent and one opponent or a multi-deal agreement consisting of multiple bilateral agreements. To clarify, we define an agreement A in the context of multi-deal as consisting of multiple *subagreements* x_1, x_2, \dots, x_k , where each subagreement x_i is drawn from outcome space Ω_i for $i = 1, 2, \dots, k$ with $k \leq m$.

A one-to-many negotiation setting does not impose restrictions on the protocol choices regarding the timing of negotiations: the center agent can negotiate with multiple opponents at the same time, often referred to as *concurrent* or sometimes *simultaneous* negotiations, or it can negotiate with opponents one after the other, called *sequential* negotiations.

Multiple negotiations can influence each other through their *interdependency*, meaning that the outcome of the complete negotiation depends on a (complex) combination of sub-agreements. For example, the fruit buyer’s smoothie is only complete when strawberries, bananas and pears are included in equal amounts. If the buyer already procured five bananas and five pears, the strategy in the subnegotiation on strawberries should be adapted. To make things more complicated, the fruit buyer’s procurements could be constrained by a limited budget. One expensive deal means there is less money available for other purchases, resulting in an updated strategy for all subnegotiations.

While traditional negotiation protocols impose strict rules, many one-to-many protocols offer agents greater flexibility to deal with these interdependencies. If an accepting message of the opponent results in an immediately binding deal for both agents, the sender agent would be forced to send out bids that are mutually exclusive. Instead, several one-to-many negotiation protocols either require an extra *confirm* message from the sender after an *accept* message from the opponent before the deal is final, or the protocols allow both agents to decommit a previous struck deal, possibly incurring a penalty.

3 Survey Methodology

With a research field as diverse as one-to-many negotiation, a careful method is required to select and analyze the most relevant research. Figure 2 illustrates how we collect, analyze and classify one-to-many negotiation papers: we review literature on one-to-many negotiation, i.e. *multiple bilateral negotiations between one agent and multiple opponents*. Other terms used for similar concepts are multi-bilateral negotiation, concurrent negotiation and multi-deal negotiation, however, their specific meaning differs between authors. The definition of one-to-many negotiation excludes bilateral negotiation, as it takes place only between one agent and one opponent, and multilateral negotiation, defined as a group of agents that shares non-bilateral communication and aims for one common deal with (a part of) the group as defined by Aydoğan *et al.* [2017]. Before, multilateral occurred as synonym for one-to-many negotiation too. We view sequential negotiation as a special case of one-to-many negotiation, provided that the agent encounters multiple opponents (bilateral repeated-encounter negotiations are not included).

In Figure 1, a visual presentation of the difference between one-to-one, one-to-many and many-to-many negotiations is shown. We consider one-to-many negotiation as a type of interaction that often occurs naturally within a many-to-many negotiation setting. For clarity, we visualize multilateral negotiations as well on the right hand side of Figure 1; however, these are not included in this survey.

This survey is an integrative literature review [Snyder, 2019] following the guidelines of Torraco [2005], specifically meant to compare research on one-to-many negotiation within the field of artificial intelligence, mathematics, computer science and alike. We exclude research in, e.g., social, historical or political science. Instead, we restrict our selection of research on papers from AAI, IJCAI, AAMAS and other AI journals and conferences. We focus on recent arti-

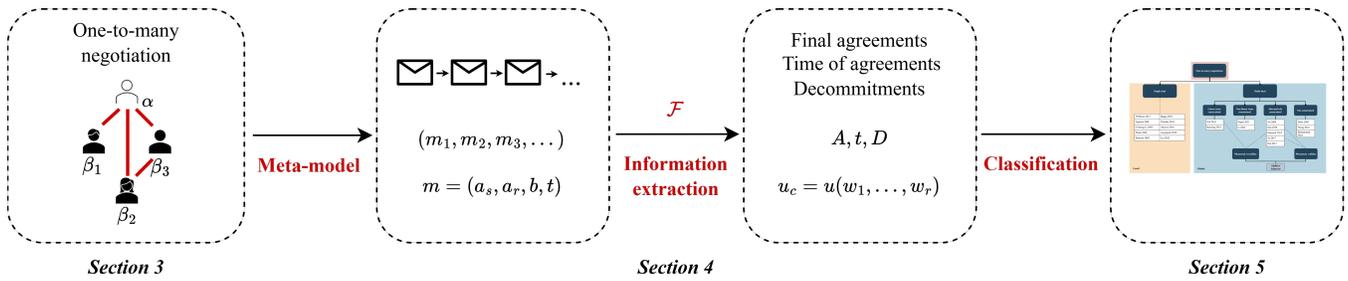


Figure 2: A visual representation of our method to quantitatively analyze one-to-many negotiations.

cles, i.e., from the years 2015 to 2024, with the exception of a few foundational classics in one-to-many negotiation that appeared before these dates.

We perform a Google Scholar search to find research articles with the following terms in the title: one-to-many negotiation (10 hits), concurrent negotiation (31) and multiple negotiations (215). Furthermore, we used “multi(-)lateral negotiation” (230 + 6) as a search term, since these terms have been used interchangeably in the past. We extend the search of the word “negotiation” by also looking for “negotiations”, “negotiate”, and “negotiating”. In total, we survey 22 articles.

4 One-To-Many Negotiation Meta-Model

To compare different one-to-many approaches, we propose a meta-model that characterizes one-to-many negotiations as a series of messages between agents (e.g., making offers, walking away). The meta-model allows us to extract all relevant information to determine the outcome of the negotiation, resulting in Table 1. We then use this information to classify the papers by *how* the center agent uses this information to compute the final utility.

We will exemplify our meta-model and information extraction with a running example of the classic paper on one-to-many negotiation by [Williams *et al.*, 2012]. They study a setting where multiple concurrent bilateral negotiations take place on a set of outcomes Ω , where $\Omega = V_1 \times V_2 \times \dots \times V_l$ denotes a contract space consisting of l negotiation issues. The center agent aims for a single deal it can locate in limited time, as the utility of a bid is discounted over time t . For flexibility, the protocol allows to decommit from previous commitments against a cost c and adds a *confirm* message for a deal to become binding.

4.1 Negotiation Messages

In our meta-model, we generalize over specific protocol choices by regarding a one-to-many negotiation as a sequence N of z general messages of the form

$$N = (m_1, m_2, \dots, m_z).$$

The protocol decides on the semantics of these messages; for example, we can replicate the setting of Williams *et al.* [2012] with a message defined as

$$m = (a_s, a_r, b, t),$$

where a_s is the agent sender, a_r is the agent receiver, b is the body of the message, and t is the time of sending. The body of the message is one of the following:

$$\{o, \text{accept}, \text{confirm}, \text{decommit}, \text{end}\},$$

where o denotes the content of the offer (i.e., the specifics of a contract; generally a mapping from negotiation issues to values). While the specific terminology can vary (e.g., using *acknowledge* instead of *confirm*, or *quit* instead of *end*), almost all papers that we survey here allow the body of the message to be (a subset of) the above. Only Alrayes *et al.* [2018] and Bagga *et al.* [2021] include special messages such as reservation requests to anticipate decommitment. Some others include additional communication, such as cheap talk [Byde *et al.*, 2003] or pre-bargaining in supply chain management [Wong and Fang, 2010].

Protocols can also impose limitations on the order in which messages can occur; for instance, a decommitment message may only follow a confirm message. Note that the order is not necessarily sequential: the same timestamp t can occur multiple times in N , for instance in synchronous protocols. The time t may even be empty for protocols where time does not play a role in the final utility. When classifying papers in this survey, rather than describing the specifics of each protocol, we will assume that the sequence of messages adheres to it at all times.

4.2 Information Extraction

In principle, the sequence of messages N consists of all the information necessary to determine the outcome utility $U : N \rightarrow \mathbb{R}$ of the negotiation. In almost all cases, however, the utility depends on only a fraction of the information contained in N : for instance the deals that were struck, but additional aspects may come to bear, such as the time of agreements and decommitments. For this, we introduce an information extractor function \mathcal{F} that takes the sequence of messages as input and produces the *utility input*; i.e., the information necessary to calculate the final utility.

In the case of [Williams *et al.*, 2012], the utility input consists of:

$A : V_1 \times V_2 \times \dots \times V_l = \Omega$: The final agreement,

$|D| = d : \mathbb{N}$: The number of decommitments,

$t : \mathbb{N}$: The time of agreement.

<i>Output \mathcal{F}: Utility input</i>	<i>Corresponding papers</i>
A: Final agreements	[Rahwan <i>et al.</i> , 2002; Byde <i>et al.</i> , 2003; Pinho <i>et al.</i> , 2004; Nguyen and Jennings, 2005; Cuihong Li <i>et al.</i> , 2005; An <i>et al.</i> , 2008; Wong and Fang, 2010; Sim and Shi, 2010; Williams <i>et al.</i> , 2012; Mansour and Kowalczyk, 2015; Yu <i>et al.</i> , 2017; Niu <i>et al.</i> , 2017; Thomas, 2018; Alrayes <i>et al.</i> , 2018; Niu <i>et al.</i> , 2018; Al-Jaljoui <i>et al.</i> , 2018; Najjar <i>et al.</i> , 2021; Mohammad, 2021; Bagga <i>et al.</i> , 2021; Baarslag <i>et al.</i> , 2021; Liu <i>et al.</i> , 2022; Li <i>et al.</i> , 2024]
t : Time of agreement	[Byde <i>et al.</i> , 2003; Williams <i>et al.</i> , 2012; Thomas, 2018; Bagga <i>et al.</i> , 2021]
D: Decommitments, including time of decommitment	[Pinho <i>et al.</i> , 2004; An <i>et al.</i> , 2008; Sim and Shi, 2010; Williams <i>et al.</i> , 2012; Yu <i>et al.</i> , 2017; Alrayes <i>et al.</i> , 2018; Bagga <i>et al.</i> , 2021]

Table 1: The utility input of the surveyed papers.

We can now formulate the utility function u_c of the center agent in terms of the input utility as follows:

$$u_c(A, d, t) = u(A) \cdot f(t) - g(d),$$

where $u(A) : \Omega \rightarrow \mathbb{R}$, $f(t) = (\frac{t}{t_{\max}})^\delta$, and $g(d) = c \cdot d$, with c a constant real value.

Existing work on one-to-many negotiation uses many different types of center utility functions; some allow for only one deal (such as the one above), some maximize or sum up the utility over all deals, while others combine them using complex, non-linear constraints. We performed the same analysis as above for each paper in our survey, identifying the template for the sequence of messages, extracting the utility input (see Table 1), and identifying the center utility function used. The resulting unified format of messages and utility functions allow us to analyze the survey papers in an integrative manner.

5 Taxonomy of One-To-Many Negotiations

Using the meta-model, we systematically collect information on the surveyed papers, which allows us to perform a structural analysis on the utility function. We introduce our taxonomy and classification for the surveyed papers from the perspective of an agent designer, focusing mainly on the interplay between the utility function and the final agreements, as that is the most influential factor on agents' strategies. Decommitments and time discount will not be our main focus, because their effects on the final utility are typically minimal. This taxonomy can guide the agent designer by (1) identifying properties of the specific problem of negotiation that they are dealing with and (2) indicating closely-related literature that discusses similar types of problems as can be obtained from the classification in Figure 3.

5.1 Number of Deals

When two agents negotiate, they commonly agree on just one final agreement (an exception is repeated encounters [Shen *et al.*, 2002]). In contrast, most one-to-many negotiations allow the agreement to consist of multiple subagreements, since the center agent can interact with multiple opponents concurrently or sequentially. Multi-deal negotiation has challenges in protocols and strategies typical of this type of negotiation because different subnegotiations influence each other [Baarslag, 2024; Florijn, 2024].

To clearly distinguish between single and multi-deal, we inspect the cardinality of final agreements A that function \mathcal{F} extracts from the negotiation. If $|A| = 1$, it is classified as single deal; if $|A| > 1$, it is classified as multi-deal. That means that a negotiation that requires one final deal [Williams *et al.*, 2012] is single deal, even if decommitment is allowed.

Intuitively, one might regard single deal negotiation as uncomplicated because of the absence of interdependency. However, the surveyed single deal papers, forming nearly 50% of the total, show subtle interdependency, which varies depending on the utility function. For example, Williams *et al.* [2012] and Alrayes *et al.* [2018] both discuss single deal protocols, where their strategies take interdependency into account in different ways, including increasing the target utility depending on the current established agreement. Multi-deal negotiations, however, often have more complex utility functions and thus show stronger interdependencies.

To accommodate interdependencies in one-to-many negotiations, approximately 40% of the surveyed protocols are specifically adapted by introducing decommitments [Nguyen and Jennings, 2005; Williams *et al.*, 2012; Alrayes *et al.*, 2018], and one third of the protocols include *confirm* messages for a deal to become binding [Alrayes *et al.*, 2018; Baarslag *et al.*, 2021; Bagga *et al.*, 2021]. This allows negotiation agents to make use of the interdependencies and take risks by using the option to quit a negotiation. However, these differences make it hard to use an agent specifically designed for one protocol in another protocol without adaptations.

An interesting side case is multi-deal bilateral negotiation, where a center agent can consolidate multiple deals with *each* opponent within one subnegotiation, and in the end select the most beneficial combination of subagreements. This combination can consist of multiple deals from one opponent as well. To our knowledge, [Yu *et al.*, 2017] is one of the few works analyzing multi-deal bilateral negotiation.

5.2 Locality of the Utility Function

In one-to-many negotiations, it is common that a coordinator keeps track of the final goal by aligning the strategies in subnegotiations. The coordinator could give instructions concerning the utilities of the potential deals in the subnegotiations, usually with a target utility [Rahwan *et al.*, 2002; Cuihong Li *et al.*, 2005; Williams *et al.*, 2012]. This is only possible if the utility function for each subnegotiation is well-defined, i.e., the center utility function is a function of the util-

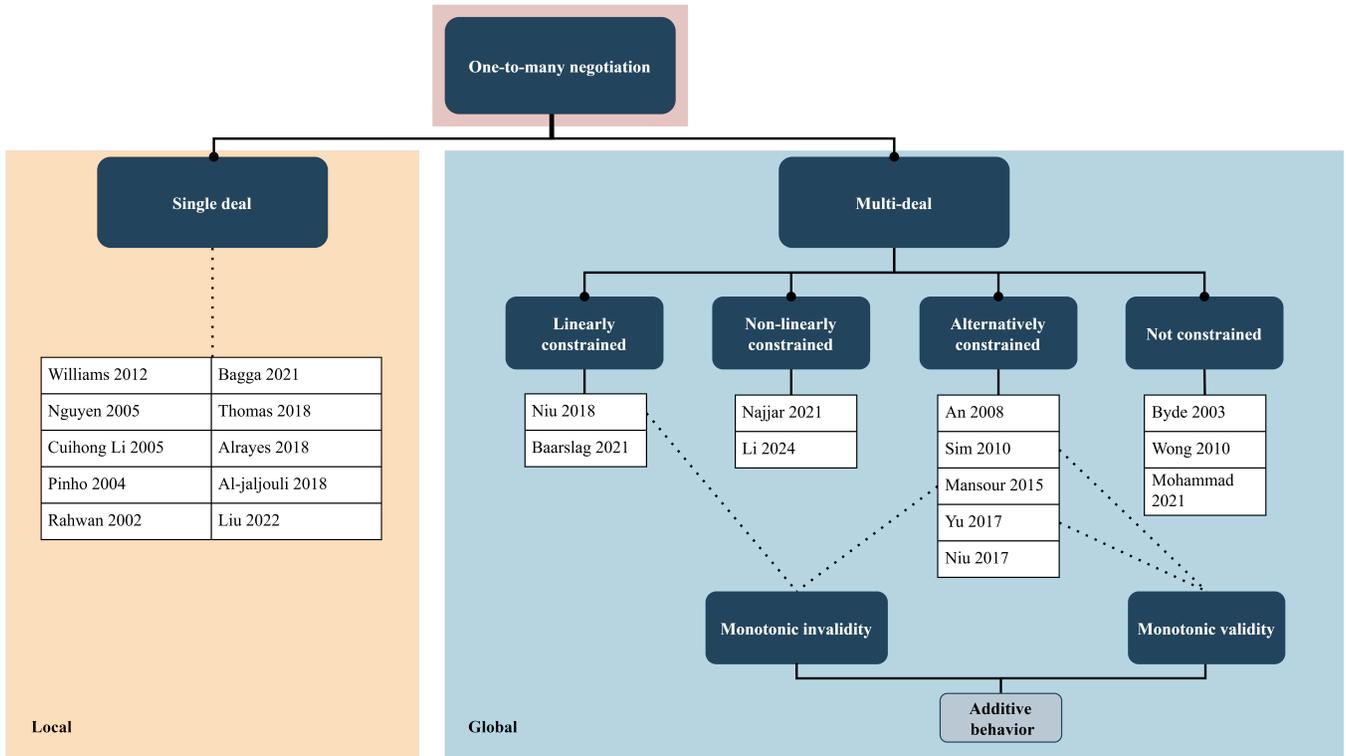


Figure 3: This figure visualizes our taxonomy. The yellow (left) and blue (right) box represent the classes *local* and *global*, respectively.

ities of each subagreement. We call these *local* utility functions. That is, for m opponents and outcomes x_1 to x_m :

$$u_c(x_1, x_2, \dots, x_m) = f(u_1(x_1), u_2(x_2), \dots, u_m(x_m)).$$

Note here, x_i could be a non-agreement (either a failed or a decommitted deal), denoted by $x_i = \emptyset$.

For example, Williams *et al.* [2012] (see Section 4.2) defines an additive function as the center utility function, so the final center utility is a summation of the utilities achieved in the subnegotiations, which is local. A local function often demands a strategy to maximize the utility in the subnegotiation. The targets in the subnegotiations, which may change over time, are based on the utility of the agreements, not the specific values (e.g., the exact number of pears). Since the strategies in the subnegotiations are relatively independent and thus simpler, one can reuse existing bilateral strategies that have been researched thoroughly [Fujita *et al.*, 2016; Baarslag *et al.*, 2014].

In contrast to local utility functions, *global utility functions* as introduced in [Mohammad, 2021] are undefined on the level of subagreement utility. In these cases, the center utility function is only defined for complete agreements, so there is no defined utility function for each subagreement with one opponent (unless one assumes disagreement for all other deals). A small change or addition of one subagreement could have a large influence and lead to a completely different outcome. Therefore, these utility functions pose a hard challenge as they do not allow a simple reuse of independent bilateral negotiation strategies in the subnegotiations.

By incorporating an arbitrary global utility function as in [Wong and Fang, 2010; Mohammad, 2021], any subagreement potentially exerts any influence on the final outcome. However, most global utility functions exhibit additional structure that the agent designers can use to their advantage. For example, Mansour and Kowalczyk [2015] discuss a buyer-seller scenario with budget constraints, requiring a global definition of utility; however, on a local level a buyer would prefer a lower over a higher price for the same object. 75% of the survey papers that use global utility functions use a type of additional constraints. This is in line with the aim of most negotiation research to mirror real world problems, as most real life applications exhibit additional structure in preferences.

Comparing global and local functions in Figure 3 may give the impression that global and local functions coincide with single deal and multi-deal, respectively; however, this is not necessarily the case. While all global utility functions are multi-deal, local utility functions can be either single deal or multi-deal. An example of a local multi-deal setting is an adaptation of [Mansour and Kowalczyk, 2015] where an agent receives the maximum value from all identified deals, but without the originally imposed budget constraints.

5.3 Validity Constraints

When negotiating multiple deals, agents often need to adhere to requirements such as achieving a minimum number of deals, maintaining a designated budget, or meeting specific targets. These constraints are imposed on the final agreement, where only some combinations of subagreements are

valid. This class is referred to as *validity-constrained* negotiations, which is a subclass of multi-deal negotiations. All constrained utility functions are global functions, as the final outcome is determined by the constraints, requiring adaptive coordination strategies to account for this global scope.

Examples of validity constraints include: “There should be a deal with all opponents” [Mansour and Kowalczyk, 2015] or “A goal should be reached with a certain tolerance to variance” [Baarslag *et al.*, 2021]. Such constraints can be expressed using a conditional statement that excludes a part of the outcome space, that is, the utility of an agreement A not in the set of valid agreements $V \subseteq \Omega$ would always be zero:

$$u_c(A) = \begin{cases} u(A) & \text{if } A \in V, \\ 0 & \text{otherwise.} \end{cases} \quad (1)$$

When striking multiple deals, the center agent is often constrained in a particular way. Firstly, *linearly constrained* utility functions show linear constraints between issues of different subnegotiations. An example of this type is a constraint of a total budget [Niu *et al.*, 2018], or a constraint stipulating that the price and quantity of goods fall within a tolerance interval of the target [Baarslag *et al.*, 2021]. Even though linear constraints are often researched in other areas such as operations research or supply chain management, less than 20% of multi-deal negotiations use only linear constraints, giving room for collaboration. Constraints of similar complexity can be expressed using a specific utility function as well [Byde *et al.*, 2003; An *et al.*, 2008; Najjar *et al.*, 2021], which simply assign lower utilities to offers outside the budget using for example a quadratic function; these fall outside this class. Secondly, *non-linearly constrained* utility functions show more complex interdependencies between issues [Najjar *et al.*, 2021; Li *et al.*, 2024]. For example, suppose that a fruit buyer would like to make a fruit salad, allowing the number of bananas and pears to show a ratio between 1:1 and 1:2. Non-linear constraints describe such proportional relations.

However, not all dependencies can be expressed using linear or non-linear constraints on issues. This third type is *alternatively constrained* utility functions. For example, in Mansour and Kowalczyk [2015] and Yu *et al.* [2017], the agent aims to acquire a certain set of objects, receiving nothing if not all are found. 80% of the alternatively constrained papers impose a similar target constraint to collect a number of items. Since almost all papers in this category allow for decommitment, the buyer has the advantage of first aiming to find a valid combination (the target items) and then continue for more efficient deals. To compensate for the buyer’s privileged position, all impose a decommitment penalty [An *et al.*, 2008; Sim and Shi, 2010; Yu *et al.*, 2017].

Finally, not all papers impose constraints, which are mostly (two out of three) in the field of supply chain management. This field commonly addresses far-reaching problems, such as the production planning, inventory management, transportation and communication with suppliers [Wong and Fang, 2010]. Even the “simpler” challenge of one-to-many negotiation has not been solved efficiently, let alone more

complex problems that involve additional complicating factors. Therefore, applied fields would benefit from extending negotiation incrementally to applicable problem settings.

5.4 Additive Behavior

Some utility functions show a particular pattern when a sub-agreement is added to the final agreement. For example, the utility of a buyer collecting goods increases when a new sub-agreement is struck as long as the budget has not been exceeded. Once the budget is exceeded, adding more bids will never help. This pattern allows the agent to search through the outcome space and direct its strategies in subnegotiations, for example in choosing to decommit existing agreements or finding new agreements to find a valid agreement. We distinguish two different types of *additive utility behavior*. A *monotonic invalidity* utility function prevents an agent from continuously seeking additional subagreements, as an invalid agreement can never turn valid by adding another subagreement. An example of this type is a budget constraint, as shown also in Mansour and Kowalczyk [2015], once a budget is exceeded, adding subagreements does not resolve the issue of the budget overrun. Formally, this class considers utility functions such that for any *invalid* agreement A , if $B \supset A$, then B is also invalid, that is,

$$\text{if } A \notin V \text{ then } B \notin V.$$

The pattern of a monotonic invalidity function can help the agent finding an endpoint in adding subagreements, which reduces the search space and thus simplifies the negotiation [An *et al.*, 2008].

A *monotonic validity* utility function is a function that preserves validity, so a valid agreement continues to remain valid when a subagreement is added to it. An example of this type of utility function is seen in the setting of Yu *et al.* [2017], where an agent needs to collect a number of goods while being allowed to freely decommit from previous subagreements. Formally, this class considers functions such that for any *valid* agreement A , if $B \supset A$, then B is also valid, that is,

$$\text{if } A \in V \text{ then } B \in V.$$

In this class, a center agent holds an advantageous position: it can first search for a satisfactory set of deals and then continue looking for even better ones [Sim and Shi, 2010; Yu *et al.*, 2017]. All surveyed papers of this type allow decommitments, putting the buyer in an even stronger and beneficial position. This is different from similar buyer-seller scenarios where agents collect particular objects but are not allowed to decommit: the center agent cannot perform optimization techniques after the deals are found [Mansour and Kowalczyk, 2015].

5.5 Message Space

Intuitively, the difficulty of a negotiation task increases as the number of actions an agent can choose from grows. For example, a scenario with many outcomes (e.g. various fruits at different prices) compared to a scenario with only a few options (e.g. one banana or one pear at a constant price), the

task becomes more complex for the agent, also in terms of computational load.

The meta-model allows us to use the size of the message space to give insight into the complexity of negotiations with respect to the domain size and the protocol. The message space dimension does not only depend on the domain size, but also on the types of agents' actions allowed and the complexity of the protocol. For example, if opponents can dynamically appear, as in [Cuihong Li *et al.*, 2005], this adds to the number of possible actions and situations that an agent would need to consider. The size of the message space captures these factors, which may influence the difficulty of the negotiation challenge.

However, it is important to note that a larger message space does not necessarily lead to a more difficult negotiation challenge. For instance, if an agent has the flexibility to decommit freely from previously made agreements, the number of possible actions increases compared to a scenario where decommitment is not permitted. However, free decommitment simplifies the negotiation challenge as all deals are non-binding. Consequently, when comparing different negotiations based on the message space, it is necessary to approach the analysis with nuance.

6 Conclusion & Discussion

In this paper, we present a taxonomy of one-to-many negotiations based on the interdependencies between them. The unified view we adopt in this survey helps to relate hitherto unconnected strands of negotiation research and to guide a negotiation agent designer in selecting the appropriate techniques for their setting. In this section, we share insights gained in designing our taxonomy and discuss aspects beyond the current state-of-the-art, followed by future directions.

Using our meta-model, we are able to compare and classify work that is based on a wide variety of protocols; however, the interpretation of the protocol and exact definition of the negotiation domain can affect the classification of the research in subtle ways. For instance, the setting of Mansour and Kowalczyk [2015] requires the center agent to find exactly one deal with every opponent; since this requirement cannot be expressed using (non-)linear issue constraints, we classify the utility function as alternatively constrained. It is trivial, however, to encode this requirement using linear constraints instead, simply by adding a negotiation issue to signal the opponent. Due to this type of delicate interplay, our classification exhibits some grey zones where judiciousness is warranted; when in doubt, the original protocol is the leading factor in our classification.

Our survey reveals an equal division of single-deal and multi-deal literature in one-to-many negotiation. We provide an explicit definition of multi-deal negotiation based on the number of agreements required by the center utility function, but the distinction is not always as clear-cut: for instance, many one-to-many negotiation agents are tasked to procure just one deal, but are subsequently tested in a tournament setting where they obtain multiple deals. This still classifies as (sequential) single-deal negotiation when their goal is to maximize average utility; technically, however, when the agent's

objective shifts to achieving top rank in a tournament, the aim of surpassing the score of the other participants results in interdependencies between the deals. Approaching such settings as multi-deal negotiations could guide agent designers to use similar techniques [Bagga *et al.*, 2021].

This survey takes the perspective of a designer of a self-interested, utility-maximizing center agent; however, the taxonomy could be extended to adopt more collaborative and altruistic points of view, e.g. by employing cooperative measures based on the level of preference conflict [Toyama and Ito, 2018; Kawata and Fujita, 2021] or protocol designs aiming for fairness, privacy, and feasibility [Al-Jaljoui *et al.*, 2018]. Viewed through this lens, our classification could be extended by incorporating the utility function of all parties involved in the negotiation.

Through our taxonomy, we are able to distinguish particular types of one-to-many negotiations that so far have been underexplored; we outline promising avenues of future research below.

Time discount. In real life negotiation scenarios, there is a cost associated with the time and resources required to reach an agreement. However, our overview as gleaned from our meta-model shows that from the papers we surveyed, only 25% use a form of time discount in their utility function. Therefore, we see ample scope for research on time-sensitive strategies in multi-deal negotiation, especially on how variation in time preferences between subnegotiations can influence the outcome.

Sequential multi-deal. To our knowledge, sequential negotiations that allow for multiple deals have not been researched yet. The work most closely related is Pinho *et al.* [2004], who study single deal sequential negotiation, but this does not extend to a multi-deal setting. Multi-deal sequential negotiation could function as an interesting initial research challenge in multi-deal negotiation, as it sidesteps the challenges associated with concurrent negotiation. We expect to see an uptick in research in this area, as it is the focal point of the Automated Negotiation League (ANL) of the Automated Negotiating Agent Competition (ANAC) at IJCAI 2025.

General approach. The field of one-to-many negotiation is rather diverse, boasting a wide gamut of approaches for protocol and agent design. To design generic agents for one-to-many negotiation and to facilitate their comparison, a common structure of their utility functions is essential, similar to the classes of utility functions that are studied in bilateral negotiation such as (linear) additive utility functions [Keeney and Raiffa, 1993], non-linear utility functions [Ito *et al.*, 2008], and ordinal preferences [Erlach *et al.*, 2018]. Enhancing our understanding of realistic yet tractable patterns in center utility functions would be a valuable contribution to the research field. This would provide the necessary groundwork for a general and future-proof one-to-many agent API and benchmarking suite in negotiation simulators such as Genius [Lin *et al.*, 2014] and NegMAS [Mohammad *et al.*, 2021]. Combining our taxonomy with use-cases from applied research [Mohammad, 2021; Najjar *et al.*, 2021] and small, real-life inspired examples [An *et al.*, 2008; Baarslag *et al.*, 2021; Li *et al.*, 2024] could serve as a first step in this direction.

Acknowledgments

We would like to thank Yasser Mohammad (NEC) for his helpful thoughts and our fruitful discussions. This publication is part of the Vidi project COMBINE (VI.Vidi.203.044) which is (partly) financed by the Dutch Research Council (NWO).

References

- [Al-Jaljoui *et al.*, 2018] Raja Al-Jaljoui, Jemal Abawajy, Mohammad M. Hassan, and Abdulhameed Alelaiwi. Secure Multi-Attribute One-to-Many Bilateral Negotiation Framework for E-Commerce. *IEEE Transactions on Services Computing*, 11(2):415–429, March 2018.
- [Alrayes *et al.*, 2018] Bedour Alrayes, Özgür Kafalı, and Kostas Stathis. Concurrent bilateral negotiation for open e-markets: the Conan strategy. *Knowledge and Information Systems*, 56(2):463–501, 2018.
- [An *et al.*, 2008] Bo An, Victor Lesser, and Kwang M. Sim. Decommithment in multi-resource negotiation (Short Paper). In *Proc. of the 7th International Conference on Autonomous Agents and Multiagent Systems*, pages 1553–1556, 2008.
- [An *et al.*, 2009] Bo An, Nicola Gatti, and Victor Lesser. Extending Alternating-Offers Bargaining in One-to-Many and Many-to-Many Settings. In *IEEE/WIC/ACM International Joint Conference on Web Intelligence and Intelligent Agent Technology*, volume 2, pages 423–426, September 2009.
- [Aydoğan *et al.*, 2017] Reyhan Aydoğan, David Festen, Koen V. Hindriks, and Catholijn M. Jonker. Alternating Offers Protocols for Multilateral Negotiation. In *Modern Approaches to Agent-based Complex Automated Negotiation*, volume 674, pages 153–167. Springer International Publishing, Cham, 2017. Series Title: Studies in Computational Intelligence.
- [Baarslag *et al.*, 2014] Tim Baarslag, Koen Hindriks, Mark Hendriks, Alexander Dirkzwager, and Catholijn M. Jonker. Decoupling Negotiating Agents to Explore the Space of Negotiation Strategies. In *Novel Insights in Agent-based Complex Automated Negotiation*, volume 535, pages 61–83. Springer Japan, Tokyo, 2014. Series Title: Studies in Computational Intelligence.
- [Baarslag *et al.*, 2021] Tim Baarslag, Tijmen Elfrink, Faria Mofakham, Thimjo Koça, Michael Kaisers, and Reyhan Aydoğan. Bargaining Chips: Coordinating One-to-Many Concurrent Composite Negotiations. In *IEEE/WIC/ACM International Conference on Web Intelligence*, pages 390–397, Essendon Vic, Australia, December 2021. ACM.
- [Baarslag, 2024] Tim Baarslag. Multi-deal Negotiation. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*, pages 2668–2673, 2024.
- [Bagga *et al.*, 2021] Pallavi Bagga, Nicola Paoletti, Bedour Alrayes, and Kostas Stathis. ANEGMA: an automated negotiation model for e-markets. *Autonomous Agents and Multi-Agent Systems*, 35(2):27, October 2021.
- [Brazier *et al.*, 2002] Frances M. T. Brazier, Frank Cornelissen, Rune Gustavsson, Catholijn M. Jonker, Olle Lindenberg, Bianca Polak, and Jan Treur. A multi-agent system performing one-to-many negotiation for load balancing of electricity use. *Electronic Commerce Research and Applications*, 1(2):208–224, June 2002.
- [Byde *et al.*, 2003] Andrew Byde, Mike Yearworth, Kay-Yut Chen, Claudio Bartolini, and Nir Vulkan. AutONA: a system for automated multiple 1-1 negotiation. In *EEE International Conference on E-Commerce, 2003*, pages 59–67, June 2003.
- [Cuihong Li *et al.*, 2005] Cuihong Li, K. Sycara, and J. Gimpapa. Dynamic Outside Options in Alternating-Offers Negotiations. In *Proceedings of the 38th Annual Hawaii International Conference on System Sciences*, pages 35b–35b, Big Island, HI, USA, 2005. IEEE.
- [Erlich *et al.*, 2018] Sefi Erlich, Noam Hazon, and Sarit Kraus. Negotiation Strategies for Agents with Ordinal Preferences. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*, pages 210–218, Stockholm, Sweden, July 2018. International Joint Conferences on Artificial Intelligence Organization.
- [Florijn, 2024] Tamara C P Florijn. Negotiation Strategies for Combining Partial Deals in One-To-Many Negotiations. 2024.
- [Fujita *et al.*, 2016] Katsuhide Fujita, Reyhan Aydoğan, Tim Baarslag, Takayuki Ito, and Catholijn Jonker. The Fifth Automated Negotiating Agents Competition (ANAC 2014). In *Recent Advances in Agent-based Complex Automated Negotiation*, volume 638 of *Studies in Computational Intelligence*, pages 211–224. Springer, Cham, 2016. Series Title: Studies in Computational Intelligence.
- [Ito *et al.*, 2008] Takayuki Ito, Mark Klein, and Hiromitsu Hattori. A multi-issue negotiation protocol among agents with nonlinear utility functions. *Multiagent and Grid Systems*, 4(1):67–83, May 2008.
- [Kawata and Fujita, 2021] Ryohei Kawata and Katsuhide Fujita. Cooperativeness Measure Based on the Hypervolume Indicator and Matching Method for Concurrent Negotiations. In *Advances in Automated Negotiations*, Studies in Computational Intelligence, pages 117–135, Singapore, 2021. Springer.
- [Keeney and Raiffa, 1993] Ralph L. Keeney and Howard Raiffa. *Decisions with Multiple Objectives: Preferences and Value Trade-Offs*. Cambridge University Press, July 1993.
- [Li *et al.*, 2024] Kai Li, Lei Niu, Fenghui Ren, and Xinguo Yu. An offer-generating strategy for multiple negotiations with mixed types of issues and issue interdependency. *Engineering Applications of Artificial Intelligence*, 136:108891, October 2024.
- [Lin *et al.*, 2014] Raz Lin, Sarit Kraus, Tim Baarslag, Dmytro Tykhonov, Koen Hindriks, and Catholijn M. Jonker. Genius: An Integrated Environment for Supporting the Design of Generic Automated Negotiators. *Computational Intelligence*, 30(1):48–70, 2014.

- [Liu *et al.*, 2022] Jingyi Liu, Nina Song, Sen Li, and Dong Xiao. A multilateral fair negotiation model solved by Hungarian solution. In Yingfa Lu and Changbo Cheng, editors, *International Conference on Computer Application and Information Security (ICCAIS 2021)*, page 81, Wuhan, China, May 2022. SPIE.
- [Mansour and Kowalczyk, 2015] Khalid Mansour and Ryszard Kowalczyk. An Approach to One-to-Many Concurrent Negotiation. *Group Decision and Negotiation*, 24(1):45–66, January 2015.
- [Mohammad *et al.*, 2019] Yasser Mohammad, Enrique Areyan Viqueira, Nahum Alvarez Ayerza, Amy Greenwald, Shinji Nakadai, and Satoshi Morinaga. Supply Chain Management World: A Benchmark Environment for Situated Negotiations. In *PRIMA 2019: Principles and Practice of Multi-Agent Systems*, volume 11873, pages 153–169. Springer International Publishing, Cham, 2019. Series Title: Lecture Notes in Computer Science.
- [Mohammad *et al.*, 2021] Yasser Mohammad, Shinji Nakadai, and Amy Greenwald. NegMAS: A Platform for Automated Negotiations. In *PRIMA 2020: Principles and Practice of Multi-Agent Systems*, Lecture Notes in Computer Science, pages 343–351, Cham, 2021. Springer International Publishing.
- [Mohammad, 2021] Yasser Mohammad. Concurrent local negotiations with a global utility function: a greedy approach. *Autonomous Agents and Multi-Agent Systems*, 35(2):28, June 2021.
- [Najjar *et al.*, 2021] Amro Najjar, Yazan Mualla, Kamal Deep Singh, Gauthier Picard, Davide Calvaresi, Avleen Malhi, Stéphane Galland, and Kary Främling. One-to-Many Negotiation QoE Management Mechanism for End-User Satisfaction. *IEEE Access*, 9:59231–59243, 2021.
- [Nguyen and Jennings, 2004] Thuc Duong Nguyen and Nicholas R. Jennings. Coordinating multiple concurrent negotiations. In *Proc. of the Third Int. Joint Conf. on Autonomous Agents and Multiagent Systems-Volume*, volume 3, pages 1064–1071, New York, 2004. ACM.
- [Nguyen and Jennings, 2005] Thuc Duong Nguyen and Nicholas R Jennings. Managing commitments in multiple concurrent negotiations. *Electronic Commerce Research and Applications*, 4(4):362–376, 2005. Publisher: Elsevier.
- [Niu *et al.*, 2017] Lei Niu, Fenghui Ren, Minjie Zhang, and Quan Bai. A Concurrent Multiple Negotiation Protocol Based on Colored Petri Nets. *IEEE Transactions on Cybernetics*, 47(11):3692–3705, November 2017. Conference Name: IEEE Transactions on Cybernetics.
- [Niu *et al.*, 2018] Lei Niu, Fenghui Ren, and Minjie Zhang. Feasible negotiation procedures for multiple interdependent negotiations. 2018.
- [Pinho *et al.*, 2004] Orlando Pinho, Geber Ramalho, Gustavo De Paula, and Patricia Tedesco. Sequential Bilateral Negotiation. In *Advances in Artificial Intelligence – SBIA 2004*, volume 3171, pages 526–535. Springer Berlin Heidelberg, Berlin, Heidelberg, 2004. Series Title: Lecture Notes in Computer Science.
- [Rahwan *et al.*, 2002] Iyad Rahwan, Ryszard Kowalczyk, and Ha Hai Pham. Intelligent Agents for Automated One-to-Many e-Commerce Negotiation. page 7, 2002.
- [Ruofei, 2016] Dou Ruofei. Research on the Optimization Model of the Supply Chain with Concurrent Negotiation Particle Swarm Optimization. *International Journal of u- and e- Service, Science and Technology*, 9(8):307–318, August 2016.
- [Shen *et al.*, 2002] Weiming Shen, Yangsheng Li, H Gheniwa, and Chun Wang. Adaptive negotiation for agent-based grid computing. *Journal of the American Statistical Association*, 97(457):210–214, 2002.
- [Sim and Shi, 2010] Kwang Mong Sim and Benyun Shi. Concurrent Negotiation and Coordination for Grid Resource Coallocation. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, 40(3):753–766, June 2010. Conference Name: IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics).
- [Snyder, 2019] Hannah Snyder. Literature review as a research methodology: An overview and guidelines. *Journal of Business Research*, 104:333–339, November 2019.
- [Thomas, 2018] Charles J. Thomas. An alternating-offers model of multilateral negotiations. *Journal of Economic Behavior & Organization*, 149:269–293, May 2018.
- [Torraco, 2005] Richard J. Torraco. Writing Integrative Literature Reviews: Guidelines and Examples. *Human Resource Development Review*, 4(3):356–367, September 2005. Publisher: SAGE Publications.
- [Toyama and Ito, 2018] Tatsuya Toyama and Takayuki Ito. Quantitatively Evaluating Difficulty in Reaching Agreements in Multilateral Closed Negotiation Scenarios. In *IEEE International Conference on Agents (ICA)*, pages 8–13, Singapore, July 2018. IEEE.
- [Williams *et al.*, 2012] Colin R. Williams, Valentin Robu, Enrico H. Gerding, and Nicholas R. Jennings. Negotiating Concurrently with Unknown Opponents in Complex, Real-Time Domains. In *ECAI 2012*, volume 242 of *Frontiers in Artificial Intelligence and Applications*, pages 834–839. IOS Press, Amsterdam, 2012.
- [Wong and Fang, 2010] T.N. Wong and Fang Fang. A multi-agent protocol for multilateral negotiations in supply chain management. *International Journal of Production Research*, 48(1):271–299, January 2010.
- [Yu *et al.*, 2017] Chunxia Yu, T.N. Wong, and Zhi Li. A hybrid multi-agent negotiation protocol supporting supplier selection for multiple products with synergy effect. *International Journal of Production Research*, 55(1):18–37, January 2017.