

MULTI-ATTRIBUTE AUCTIONS: APPLICATION TO WORKFLOW MANAGEMENT SYSTEMS

Albert Pla Planas

Dipòsit legal: Gi. 990-2014

<http://hdl.handle.net/10803/134731>

ADVERTIMENT. L'accés als continguts d'aquesta tesi doctoral i la seva utilització ha de respectar els drets de la persona autora. Pot ser utilitzada per a consulta o estudi personal, així com en activitats o materials d'investigació i docència en els termes establerts a l'art. 32 del Text Refós de la Llei de Propietat Intel·lectual (RDL 1/1996). Per altres utilitzacions es requereix l'autorització prèvia i expressa de la persona autora. En qualsevol cas, en la utilització dels seus continguts caldrà indicar de forma clara el nom i cognoms de la persona autora i el títol de la tesi doctoral. No s'autoritza la seva reproducció o altres formes d'explotació efectuades amb finalitats de lucre ni la seva comunicació pública des d'un lloc aliè al servei TDX. Tampoc s'autoritza la presentació del seu contingut en una finestra o marc aliè a TDX (framing). Aquesta reserva de drets afecta tant als continguts de la tesi com als seus resums i índexs.

ADVERTENCIA. El acceso a los contenidos de esta tesis doctoral y su utilización debe respetar los derechos de la persona autora. Puede ser utilizada para consulta o estudio personal, así como en actividades o materiales de investigación y docencia en los términos establecidos en el art. 32 del Texto Refundido de la Ley de Propiedad Intelectual (RDL 1/1996). Para otros usos se requiere la autorización previa y expresa de la persona autora. En cualquier caso, en la utilización de sus contenidos se deberá indicar de forma clara el nombre y apellidos de la persona autora y el título de la tesis doctoral. No se autoriza su reproducción u otras formas de explotación efectuadas con fines lucrativos ni su comunicación pública desde un sitio ajeno al servicio TDR. Tampoco se autoriza la presentación de su contenido en una ventana o marco ajeno a TDR (framing). Esta reserva de derechos afecta tanto al contenido de la tesis como a sus resúmenes e índices.

WARNING. Access to the contents of this doctoral thesis and its use must respect the rights of the author. It can be used for reference or private study, as well as research and learning activities or materials in the terms established by the 32nd article of the Spanish Consolidated Copyright Act (RDL 1/1996). Express and previous authorization of the author is required for any other uses. In any case, when using its content, full name of the author and title of the thesis must be clearly indicated. Reproduction or other forms of for profit use or public communication from outside TDX service is not allowed. Presentation of its content in a window or frame external to TDX (framing) is not authorized either. These rights affect both the content of the thesis and its abstracts and indexes.



Universitat de Girona

PhD Thesis

**Multi-attribute Auctions: Application to
Workflow Management Systems.**

Albert Pla Planas

2014



Universitat de Girona

PhD Thesis

**Multi-attribute Auctions: Application to
Workflow Management Systems.**

Albert Pla Planas

2014

Doctoral Programme in Technology

Supervised by:

Dr. Beatriz López and Dr. Javier Murillo

Work submitted to the University of Girona in partial fulfilment of the
requirements for the degree of Doctor of Philosophy

Dr. Beatriz López from Universitat de Girona.

Dr. Javier Murillo from Newronia.

DECLARE

That the work entitled *Multi-attribute Auctions: Application to Workflow Management Systems*. presented by *Albert Pla Planas* to obtain the degree in Doctor of Philosophy has been developed under our supervision and complies with the requirements needed to obtain the International Mention.

Therefore, in order to certify the aforesaid statement, we sign this document.

Girona, March 2014.

Agraïments

Dur a terme aquesta tesi no ha set fàcil. Però, per sort, he tingut molta gent al voltant que m'ho ha fet una mica menys difícil. Aquesta tesi resumeix el treball que he realitzat durant els últims 4 anys però, en el fons, no deixa de ser la cirereta d'un pastís que es va començar a cuinar quan vaig començar la carrera. Així, doncs, m'agradaria donar les gràcies a la llarga llista de persones que durant aquests anys m'han ajudat d'una manera o d'una altra i que m'han fet costat quan ho he necessitat.

M'agradaria agrair especialment a la Bea i en Javi que m'hagin guiat durant la tesi. Per ajudar-me a trobar solucions a problemes que a vegades semblaven insolucionables i veure les coses des de noves perspectives. Però, sobretot, agrair-los la paciència de no engegar-me en determinats moments.

En la mateixa línia, vull donar les gràcies a tots els companys de *La Pizzeria* i del grup de recerca. A tots els que han passat per aquí i als que hi acaben d'arribar, i especialment a en Xavi i en Ferran per les hores de pissarra, de discussions i del que hagi fet falta.

El refranyer diu *Ens ha calgut déu i ajuda arribar fins aquí*. Jo hi afegiria cafeïna. Molta cafeïna. I alguna que altra cervesa. I algun que altre sopar. I val més no seguir... Per això vull agrair aquests moments a tota la gent que he conegut durant la carrera. Menció d'honor a en Pablo, l'Enric, en Gubern, en Pons, en Ricard, en Mariano i en Massi. Probablement, sense tots aquests cafés, cerveses, viatges, nits de dijous i rareses diverses que hem viscut, el cap m'hagués esclatat en algún moment indeterminat i tot hagués set molt diferent.

A l'Anna, per estar al meu costat els dies brillants i per suportar-me els foscos (que no deu ser fàcil).

Per acabar m'agradaria donar les gràcies als meus pares per portar-me fins aquí. Sense la seva ajuda el camí no hauria estat tan llarg ni hauria valgut tant la pena.

I would like to thank all the people I've had the chance to meet during my research internship at the LIP6 research group at the UPMC-Paris6 University. Specially to Nicolas Maudet

who offered me the opportunity to visit such great university and who guided me during my stay. Also I thank the reviewers of the thesis for their helpful feedback and. Finally, I would like to thank whoever is reading this manuscript for investing his time in my work; I hope the reading is not too heavy...

Gràcies a tots.

This thesis has been developed with the support of the University of Girona (BR10/18 Scholarship) and the projects labeled TIN2008-04547 and EUREKA ITEA 2 n11027 - IPT-2012-0943-300000.

Publications

The publications derived from this PhD thesis are the following:

Journals

- **Albert Pla**, Pablo Gay, Joaquim Meléndez and Beatriz López. *Petri net-based process monitoring: a workflow management system for process modelling and monitoring*, Journal of Intelligent Manufacturing. 2012. pages 1-16 (pre-print). DOI:10.1007/s10845-012-0704-z
- **Albert Pla**, Beatriz López, Javier Murillo and Nicolas Maudet. *Multi-Attribute Auctions with Different Types of Attributes: Enacting Properties in Multi-Attribute Auctions*. Expert Systems with Applications. (In press, available online , Accepted February 4th, 2014) DOI:10.1016/j.eswa.2014.02.023.
- **Albert Pla**, Beatriz López and Javier Murillo. *Multidimensional Fairness for Auction-based Resource Allocation*. Knowledge-based Systems. (Submitted on November, 2013).

Conferences

- Beatriz López, **Albert Pla**, David Daroca, Luís Collantes, Sara Lozano and Joaquim Meléndez. *Medical equipment maintenance support with service-oriented multi-agent services*. In Principles and Practice of Multi-Agent Systems (PRIMAA). Proceedings published at Lecture Notes in Computer Science Vol. 7057. pp. 487-498. Kolkata, India, November 12th-15th, 2010.
- **Albert Pla**, Beatriz López and Javier Murillo. *Multi-Attribute Auction Mechanism for Supporting Resource Allocation in Business Process Enactment*, Sixth Starting AI Researchers'

Symposium (STAIRS 2012) ECAI 2012. Proceedings published at Frontiers in Artificial Intelligence and Applications Vol. 241. pp. 228-239 Montpellier, France, August 2012.

- **Albert Pla**, Beatriz López and Javier Murillo. *Multi Criteria Operators for Multi-attribute Auctions*. 9th International Conference on Modelling Decisions for Artificial Intelligence (MDAI 2012). Proceedings published at Lecture Notes on Computer Science Vol. 7647. pp 318-328 Girona, Spain, November 2012.
- **Albert Pla**, Beatriz López and Javier Murillo. *How to Demonstrate Incentive Compatibility in Multi-Attribute Auctions*. 16th International Conference of the Catalan Association for Artificial Intelligence. Proceedings published at Frontiers in Artificial Intelligence and Applications Vol 256. pp 303-306. Vic, Spain, October 2013.
- Ferran Torrent, **Albert Pla** and Beatriz López. *Using multi-attribute combinatorial auctions for sustainable workflow scheduling*. 28th AAAI Conference on Artificial Intelligence (AAAI-14). (Submitted on January, 2014).

Workshops

- **Albert Pla**, Beatriz López and Javier Murillo. *Workflow Resource Allocation through Auctions*. 2nd Workshop on Artificial Intelligence and Logistics IJCAI, pp 55-60, Barcelona 2010.

Side publications developed during the PhD period which are not directly linked to the thesis:

Journals

- **Albert Pla**, Beatriz López, Pablo Gay and Carles Pous. *eXiT*CBR.v2: Distributed case-based reasoning tool for medical prognosis*. Decision Support Systems Vol.54 (3), pp 1499-1510. February 2013. DOI:10.1016/j.dss.2012.12.033. ISSN: 0167-9236
- Pablo Gay, Beatriz López, **Albert Pla**, Jordi Saperas and Carles Pous. *Enabling the Use of Hereditary Information from Pedigree Tools in Medical Knowledge-based Systems*, Journal of Biomedical Informatics, Vol.46 (4), pp 710-720. August 2013. Au DOI:0.1016/j.jbi.2013.06.003. ISSN 1532-046

Workshops

- Beatriz L'opez, **Albert Pla**, Carles Pous, Pablo Gay, Joan Brunet. *System Overview: Breast Cancer Prognosis Through CBR*. Workshop on Case-Based Reasoning in the Health Sciences 2012 ICCBR2012. pp 105-110. September 2012, Lyon (France).

Acronyms and Abbreviations

BB-WOC	Bid-based won auction coefficient
BB-LS	Bid-based losing streak
BDC	Bidder drop control
BDP	Bidder drop problem
BP	Business process
CFP	Call for proposals
DPORA	Discriminatory price optimal recurring auction
GSP	Generalized Second Price
FMAAC	Framework for multi-attribute auction customization
IC	Incentive compatible
JIT	Just-in time
LS	Losing streak
MARA	Multi-attribute resource allocation
MAS	Multi-agent system
MU	Multi-unit
MURA	Multi-unit Recurrent Auctions
PERA	Preference-based english reverse auctions
POS	Probability of success
PPC	Pay per click
probBB-LS	Probabilistic bid-based losing streak
probBB-WOC	Probabilistic bid-based won auction coefficient
probLS	Probabilistic losing streak
probWOC	Probabilistic won auction coefficient
PUMAA	Preserving utility multi-attribute auction
RP	Resource provider
SA	Service agent

SFD	Spearman's footrule distance
SMT	Satisfiability modulo theories
SP	Second Price
TB	Truthful bidding
VCG	Vickrey-Clarke-Groves
VLLF-BDC	Valuable Last Loser First
WDP	Winner determination problem
WOC	Won auction coefficient
WS	Weighted sum
WSF	Weighted sum of functions
WRM	Wealth rank modification

List of Figures

2.1 Auctions described in this chapter.	10
2.2 Auction classification according to different properties.	11
2.3 Example of a Google sponsored search result.	17
2.4 Auction classification according to the WDP dimensionality	31
2.5 PUMAA shown with respect to the auctions described in Chapter 2	33
2.6 Auction style summary.	36
3.1 Workflow reference model according to [87].	39
3.2 Multi-agent system schema	40
3.3 Examples of functions that cannot be used as evaluation functions and their possible alternatives.	52
4.1 Classification of the attributes involved in a multi-attribute auction according to its origin and verifiability.	67
4.2 Example of the attribute types used in the bidding process	70
5.1 fair-PUMAA and other egalitarian auctions classification with respect to the auctions described in 2.	86
5.2 Priority methods classification.	88
6.1 Multi-agent system architecture utilized within the simulation	99
6.2 Example of the types of resource required by each service agent.	101
6.3 Proportion of each type of task in the original data.	102

6.4	Left) Probability of a task of type S3 occurring in a given time instant. Right) Probability function that defines the emergency level of a task of type S3.	103
6.5	Distribution functions that determine the duration of each type of task for resource provider RP7	104
6.6	Total task allocation cost for each service agent using different allocation mechanisms	107
6.7	Mean cost for a service allocation when using different allocation methods.	108
6.8	Delays produced in the simulation when using different allocation methods.	108
6.9	Resource providers' profits when <i>RP1</i> and <i>RP2</i> cheat (grey) and when they all bid truthfully (black).	111
6.10	Service agents expenses when <i>RP1</i> and <i>RP2</i> cheat (grey) and when they all bid truthfully (black).	112
6.11	Mean resource providers' benefits when <i>RP5</i> to <i>RP8</i> cheat (grey) and when they all bid truthfully (black).	113
6.12	Mean service agents' expenses when <i>RP5</i> to <i>RP8</i> cheat (grey) and when they all bid truthfully (black).	114
6.13	Mean service agents' utilities benefits in Experiment 2 - Scenario 2.	115
6.14	Probability for each <i>RP1</i> strategy after a 200 days simulation.	116
6.15	Evolution of the bidding strategy probabilities for the reinforcement learning agent.	117
6.16	Evaluation function impact measures	119
6.17	Resource providers benefits when using different evaluation functions.	120
6.18	Comparison between Che's second score auction and PUMAA in terms of robustness.	123
6.19	Service agents (auctioneers) mean expenditure at the end of simulations	128
6.20	Resource providers (bidders) mean revenue at the end of simulations	129
6.21	Resource agents' Gini's index for the resulting allocations in terms of revenue	130
6.22	Mean number of delays produced for the different priority methods.	131
6.23	Mean Gini Index and mean WRM for the stochastic priority methods	134
6.24	Active bidders in the market during in a BDP experiment.	137

6.25	Box plot showing the number of active bidders after the simulations	139
6.26	Box plot showing the mean cost of a workflow allocation after the simulations .	140
6.27	Mean task cost after a BDP simulation	141

List of Tables

2.1	Table of bids corresponding to Example 2.2.	16
2.2	Bids and their resulting utilities for Example 2.4	18
2.3	Resulting allocation Example 2.5 following the Google's pay per click position auction approach	19
2.4	Bids and scores for Example 2.7	22
2.5	Received bids by <i>WeMakeCars</i> and their expected utilities	27
3.1	Bids received in Example 3.3	47
4.1	Attributes involved in some of the auctions presented in this thesis.	69
4.2	List of bids and their corresponding ranks and evaluations.	77
5.1	Example of the quantitative priority methods	89
5.2	Example of the qualitative priority methods	92
6.1	Summary of the experiments performed.	98
6.2	Example of resource skills ($T =$ type, t task duration) and costs (c) randomly generated per each agent.	101
6.3	Number of instances of each type of task in the original data.	102
6.4	Relation of the tasks S that a resource of a type R can perform.	103
6.5	Relationship between the urgency of a service and its maximum execution time.	104
6.6	Example of resources skills randomly generated per each agent: (Resource type, delivery time, tolerance to errors and cost)	118

Contents

Abstract	xxi
Resum	xxiii
Resumen	xxv
1 Introduction	1
1.1 Introduction	1
1.2 Contributions	4
1.3 Thesis Outline	4
2 Auctions for M.A.R.A.	7
2.1 Introduction	7
2.1.1 Basic game theory concepts	8
2.1.2 Auction types	11
2.1.3 Auction properties	12
2.2 Uni-attribute Auctions	13
2.2.1 Vickrey auctions	14
2.2.2 Position auctions and the sponsored search market	17
2.3 Multi-Attribute Auctions	20
2.3.1 Scoring rules-based auctions	20
2.3.2 Auctions with flexible attribute structures	24

2.3.3	Auction-based allocations where the task delivery is uncertain	26
2.3.4	Attributes as constraints	29
2.4	Summary	31
2.4.1	Dimensionality	31
2.4.2	Auction properties	32
3	PUMAA	37
3.1	Assumptions and Limitations	38
3.2	Mechanism	40
3.2.1	Call for proposals	41
3.2.2	Bidding	43
3.2.3	Winner determination	45
3.2.4	Payment	47
3.3	Evaluation Function Characteristics	50
3.3.1	Multicriteria function as evaluation function: requirements	51
3.3.2	Multicriteria function as evaluation function: examples	53
3.4	Incentive Compatibility	55
3.4.1	Condition satisfaction	57
3.4.2	Proof by using a solver	59
3.5	Considerations Regarding Other PUMAA Properties	61
3.6	Summary	63
4	A Framework for Multi-Attribute Auction Customization	65
4.1	Attribute Typologies in Multi-attribute Auctions	65
4.1.1	Ownership of the attributes	66
4.1.2	Verifiability	67
4.2	FMAAC: Framework for Multi-attribute Customization	70
4.2.1	Call for proposal - Defining the bidder-provided attributes	71
4.2.2	Bidding - The role of unverifiable attributes	72

4.2.3	Winner determination - The role of auctioneer-provided attributes	73
4.2.4	Payment - Playing all the attributes together	75
4.2.5	Attribute information update	76
4.2.6	A simple example	76
4.3	Summary	79
5	Multi-dimensional Fairness for M.A.R.A.	81
5.1	Preliminary Information	83
5.1.1	Discriminatory price optimal recurring auction	83
5.1.2	Priority auctions	84
5.2	Multidimensional Fairness	85
5.3	fair-PUMAA	85
5.4	Priority Calculation Methods	88
5.4.1	Quantitative priority methods	89
5.4.2	Qualitative bid-based priorities	91
5.4.3	Stochastic priorities.	93
5.5	Considerations Regarding fair-PUMAA Properties	94
5.6	Summary	96
6	Experimentation and Results	97
6.1	Introduction	97
6.2	Simulation Environment and Data	98
6.2.1	Multi-agent simulation	98
6.2.2	Data used in the simulations	100
6.3	PUMAA Results	105
6.3.1	Experiment 1: The need of a multi-attribute auction	106
6.3.2	Experiment 2: Strategy proofness	109
6.3.3	Experiment 3: Impact of the evaluation function in the auction allocation	118
6.3.4	Experiment 4: Robustness to cheaters and utility preservation	121

6.4	Multi-dimensional Fairness Results	126
6.4.1	Experiment 1: Uni-dimensional vs. multi-dimensional fairness	126
6.4.2	Experiment 2: Stochastic priorities	132
6.4.3	Experiment 3: Tackling the bidder drop problem	136
6.5	Summary	141
7	Conclusions	143
7.1	Summary	143
7.2	Future Work	145
A	Notation Guide	149
A.1	PUMAA	149
A.2	FMAAC	151
A.3	Multi-dimensional Fairness	153
	Bibliography	155

Abstract

Resource and task allocation for workflows poses an allocation problem in which several attributes may be involved (economic cost, delivery time, CO₂ emissions...), therefore, it must be treated from a multi-criteria perspective so that all of the attributes are taken into account when deciding the optimal assignments. Auction mechanisms offer the chance to allocate resources and services in a competitive market environment whilst optimizing outcomes for all of the participants. In this thesis, we propose the use of multi-attribute auctions for allocating resources to workflows occurring in dynamic environments where task performance is uncertain.

For that purpose, we present PUMAA, a multi-attribute auction mechanism specially designed for allocating tasks in workflows and to reduce auctioneer's utility loss in case of task misdelivery. PUMAA follows a Vickrey auction structure and, conversely to previous approaches, uses a conditional payment method which relates the payment to the quality of the delivered task. Moreover, PUMAA ensures incentive compatibility (the best strategy for bidders is truthful bidding), a desired property for guarantee mechanism outcomes and elude misbehavior. This property is mathematically demonstrated and experimentally corroborated.

Next, we deepen on the different types of attributes which can take part into an auction and we classify them according to their origin and their verifiability. Using this classification, we present a framework for multi-attribute auction customization (FMAAC). FMAAC allows mechanism designers to include new properties (such as fairness, robustness or reliability) to the auction mechanism depending on the type of attributes they are using.

Following FMAAC, the thesis focusses on fairness issues in auctions. Multi-attribute resource allocation problems involve the allocation of resources on the basis of several attributes, therefore, the definition of a fairness method for this kind of auctions must be formulated from a multi-dimensional perspective, as focusing in just a single attribute may compromise the allocations regarding the remainder. In particular, we present a multi-dimensional fairness approach based on priorities that we call fair-PUMAA, a multi-attribute auction mechanism

which favors egalitarian allocations.

The methodology described in this manuscript is tested in a multi-agent system-based simulation framework which emulates the behavior of a manufacturing industry. Results support the statements that PUMAA is incentive compatible and that it preserves the auctioneer's utility in case of unsuccessful tasks. Regarding fair-PUMAA, experimentation shows that the use of multi-dimensional fairness produces egalitarian allocations that can be useful to alleviate issues such as the bidder drop problem.

Resum

L'assignació de tasques i recursos en fluxos de treball planteja un problema en el qual poden intervenir diferents atributs (costs econòmics, terminis d'entrega, emissions de CO₂, etc.). Conseqüentment, per tal de tenir en compte tots els elements involucrats en l'assignació de recursos i aconseguir-ne una d'òptima, cal enfocar el problema des d'un prisma multicritèria. Els mecanismes de subhasta ofereixen la possibilitat d'assignar recursos i serveis en entorns competitius al mateix temps que s'optimitzen els beneficis per a tots els participants. En aquesta tesi es proposa emprar subhastes multi-atribut per a l'assignació de recursos en fluxos de treball que es desenvolupen, concurrentment, en entorns dinàmics on el desenvolupament de tasques presenta un alt grau d'incertesa.

Amb aquest objectiu, es presenta PUMAA: un mecanisme de subhastes multi-atribut dissenyat específicament per assignar recursos en fluxos de treball i per reduir la pèrdua d'utilitat que els subhastadors poden patir en cas d'incompliments en l'execució de tasques. PUMAA segueix l'estructura d'una subhasta de Vickrey però, en contraposició a mètodes anteriors, utilitza un mètode de pagament condicional que lliga el preu pagat per una tasca amb la qualitat en la qual aquesta s'ha desenvolupat. A més, PUMAA incentiva als licitants a revelar les seves preferències quant a preu i els altres atributs ("Incentive Compatibility"), una propietat essencial per garantir la maximització dels beneficis dels diferents participants de la subhasta i evitar l'aparició de postors amb comportaments maliciosos. Aquesta propietat es demostra matemàticament i experimentalment.

En el transcurs d'aquesta tesi s'estudien els diferents atributs que poden prendre part dins d'una subhasta i es classifiquen segons el seu origen i la seva verificabilitat. Mitjançant aquesta classificació, es desenvolupa un entorn per a la personalització de subhastes multi-atribut (FMAAC). Mitjançant FMAAC, els dissenyadors de subhastes poden incorporar diferents propietats (per exemple: equitat, robustesa o fiabilitat) en les subhastes segons les particularitats del domini on es desenvolupen i els atributs que s'emprin.

Els problemes d'equitat que poden derivar de les subhastes multi-atribut també són estudiats

en aquesta dissertació. La dimensionalitat present en l'assignació de recursos afecta també la formulació de mètodes d'equitat: focalitzar l'equitat en un determinat atribut podria posar en perill la qualitat de les assignacions respecte a la resta d'atributs. Conseqüentment, es desenvolupa un mètode d'equitat multi-dimensional basat en prioritats anomenat fair-PUMAA: un mecanisme de subhastes multi-atribut que afavoreix les assignacions igualitàries.

La metodologia presentada al llarg d'aquesta tesi es contrasta a través d'un entorn de simulació multi-agent que emula el funcionament d'una indústria on es desenvolupen diferents processos industrials. Els experiments corroboren les hipòtesis que PUMAA incentiva als postors a revelar els seus valors veritables i que PUMAA minimitza la pèrdua d'utilitat dels subhastadors en cas que les tasques no es desenvolupin com s'havia acordat. Els diferents experiments també indiquen que fair-PUMAA afavoreix la producció d'assignacions igualitàries i que els mètodes d'equitat multidimensionals poden moderar problemes com la pèrdua de licitadors en subhastes recurrents (Bidder drop problem).

Resumen

La asignación de tareas y recursos en flujos de trabajo plantea un problema en el cual pueden intervenir distintos atributos (costos económicos, plazos de entrega, emisiones de CO₂, etc.). En consecuencia, para tener en cuenta todos los elementos involucrados en el proceso y para obtener una asignación óptima es necesario plantear el problema con un enfoque multicriteria. Los mecanismos de subasta ofrecen la posibilidad de asignar recursos y servicios en entornos competitivos al mismo tiempo que se optimizan los beneficios de todos los participantes. En esta tesis se propone usar subastas multi-atributo para la asignación de recursos en flujos de trabajos que se desarrollan concurrentemente en entornos dinámicos donde la ejecución de tareas presenta un alto grado de incertidumbre.

Con este objetivo, se presenta PUMAA: un mecanismo de subastas multi-atributo diseñado específicamente para asignar recursos en procesos de producción y para reducir la pérdida de utilidad que los subastadores pueden sufrir en caso de incumplimientos en la ejecución de tareas. PUMAA sigue la estructura de una subasta de Vickrey pero, a diferencia de otros métodos, utiliza un mecanismo de pago condicional que relaciona la cantidad a pagar por el desarrollo de una tarea con la calidad con la que ésta se ha desarrollado. Además, PUMAA incentiva a los licitadores a revelar sus preferencias reales en lo que se refiere a precio y los demás atributos ("Incentive Compatibility"), una propiedad esencial para garantizar la maximización de los beneficios de los participantes y para evitar la aparición de licitadores con comportamientos maliciosos. Esta propiedad se demuestra tanto matemáticamente como empíricamente.

En el transcurso de esta tesis se analizan los distintos atributos que pueden usarse en un proceso de subasta y se clasifican en base a su origen y a su verificabilidad. Mediante dicha clasificación, se desarrolla un entorno de trabajo para la personalización de subastas multi-atributo (FMAAC). Usando FMAAC y escogiendo los atributos adecuados, los diseñadores de mecanismos de subasta pueden incorporar a las subastas distintas propiedades (por ejemplo equidad, robusteza o fiabilidad) acorde con las particularidades del dominio donde se desarrolla la actividad.

Los problemas de equidad que pueden derivar de las subastas multi-atributo también se estudian en este documento. La dimensionalidad presente en la distribución de recursos afecta la formulación de métodos de equidad: centrar la equidad en un solo atributo podría comprometer la calidad de la asignación en relación a los demás atributos. En consecuencia, se diseña fair-PUMAA: un método de equidad multidimensional para subastas basado en prioridades que favorece una distribución de recursos más igualitaria.

La metodología presentada a lo largo de esta tesis se contrasta mediante un entorno de simulación multi-agente que emula el desarrollo de una industria donde se suceden distintos procesos productivos. Los experimentos corroboran las hipótesis planteadas: PUMAA incentiva los licitadores a revelar sus verdaderos valores y minimiza la pérdida de utilidad de los subastadores cuando se producen fallos en la ejecución de las distintas tareas; por otra parte, los experimentos sugieren que fair-PUMAA favorece la distribución igualitaria de los recursos y que los métodos de equidad multi-dimensional pueden minimizar problemas como la pérdida de licitadores en subastas recurrentes (Bidder drop problem).

INTRODUCTION

This chapter will provide an overview of this PhD thesis and presents the scope, the hypothesis and the contributions made in this dissertation. Finally, an outline of the contents of each chapter is described.

1.1 Introduction

Business process management is becoming a fundamental element of many industrial processes. In today's economy, suppliers, manufacturers and retailers are working together to reduce production costs, maximize productivity and to optimize product quality and production times. To manage the evolution and interactions of different business actions it is necessary to accurately model the various steps in the process concerned, the resources needed and the flow of messages between the different parties involved (suppliers, manufacturers, clients, etc.). Workflows provide a way of describing the order of execution and the dependence relationships between the constituent activities of business processes [29, 1].

Workflows usually model single and unique business processes, although in real life environments processes represented in workflows are rarely executed in isolation [82]. Workflows are usually executed concurrently, using a limited number of local resources but also outsourcing tasks to external providers. In consequence, a delay in an ongoing workflow can impact other pending workflows. This causes a cascade effect in the performance of the rest of the system, due to dependencies on resources and occupations. For this reason it is important to monitor not only a single workflow's execution but the entire system as a whole. As a delay can disrupt in the execution of other tasks [86]. Moreover, these dependencies enhance the need to obtain efficient allocations to external providers not only in terms of revenue but also in terms of

delivery time and quality preservation. Thus, the methodology used to assign tasks to external providers becomes a crucial issue for the manufacturing and supply-chain industries. Auction mechanisms offer the ability to allocate resources and services in a market (e.g. a company that desires to externalize a production task [67, 78]) while optimizing the outcomes for all of the participants (both buyers and sellers).

Auctions are gaining relevance in today's economy. They are increasingly being used in fields such as procurement [77], supply chain management [71], electronic advertisement [28], the smart grid [66], service allocation [79] and electronic commerce [74]. For instance, using auctions, manufacturers can obtain feedstock from their providers under the best economic conditions and popular web sites can sell their advertising space to the marketing companies which offer the most profitable revenue. In an auction, the auctioneer (e.g. manufacturer) that wishes to buy a product announces this requirement in a call for proposals. Interested agents, bid for this at a price they consider suitable to deliver the service and make an acceptable profit. The auctioneer can clear the auction, by determining a winner, and then setting up the price of the goods.

For example, in certain auctions, all the bids that do not achieve a minimum quality level are filtered out and discarded before determining the winner. Multi-attribute auctions have been designed to ensure multiple attributes are considered in the winner determination problem. In this case, the auctioneer is faced with the problem of choosing a winner among a set of Pareto-optimal solutions. For example, in a multi-attribute auction in which an auctioneer wishes to externalize a task, the auctioneer can characterize the task according to different attributes, such as price and quality. The auctioneer desires that task is performed at the highest quality and at a good price. He will receive different bids of different qualities and prices. Quality and price are often conflicting criteria and several attribute combinations may be equally optimal. Thus, the auctioneer should make a decision and select a combination of attributes. Depending upon how this selection is performed, truthful bidding based exclusively on price is no longer valid. For example a bidder which is encouraged to provide its true value for the price, may provide a fake proposal concerning quality, resulting in harmful consequences for the auctioneer. In such scenarios, where money is not the sole issue (e.g. pay-per-click advertising), new auction mechanisms have been designed. An example of this are position auctions used in advertisement scenarios such as Google Ads [84] where auctions occur in a repetitive manner amongst a specific set of bidders. Due to repetition, auctioneers are able to qualify bidders according to various attributes. When clearing the auction, the auctioneer unifies the qualification attributes and the received bids in order to determine the winner. By doing so, this kind of auction guarantees incentive compatibility, as the attributes this mechanism deals

with are not manipulable by bidders, as they are provided by the auctioneer itself.

In this dissertation we distinguish different types of attributes regarding the goods or product to be sold, verifiable and unverifiable attributes, in addition to attributes qualifying the buyers. Verifiable attributes are the ones that can be checked upon the receipt of the goods (e.g. delivery time or physical properties). Conversely, unverifiable attributes, cannot be checked due to their subjective nature. The price of the object (a good can be expensive or not depending on the utility provided to the buyer) is an example of this kind of attribute. The third kind of attribute, auctioneer provided attributes, concerns information about goods itself and/or the bidder itself (e.g. the reliability of a bidder).

This classification allows the creation of a new auction mechanism that takes into account results of previous uni-attribute and multi-attribute auctions. This system allows the auctioneers to acquire some desirable properties such as incentive compatibility. In this way bidders are encouraged to bid based upon true properties and attributes they can provide (duration, quality of service) at the best price. However, the use of different types of attribute also allows to alter the structure of the auction in order to customize the resulting allocations. For instance, obtaining a fairer allocation (an allocation which all the participants consider as reasonable so they continue to participate in the system); or obtaining more robust allocations when the commitment of the bidder to the tasks can not be guaranteed (the auctioneer will assign the task to those agents most likely to finish their jobs in an appropriate manner).

This thesis covers a wide scope of domains. The main goal of this research is to provide a mechanism for allocating tasks in uncertain production domains, in an efficient manner. This is in terms of cost but also in terms of important attributes for the supply chain domain, such as delivery times, delay reduction or quality. For this purpose we first present a literature review concerning auctions from a workflow management perspective. Second, we define a multi-attribute auction mechanism for task and resource allocation in uncertain production environments. This allows us to combine several attributes (such as starting time, price, quality, etc.), whilst incentivizing the bidders to bid truthfully, ensuring effective allocations. Then, by classifying the attributes involved in an auction, we propose a framework for extending such a mechanism in order to allow the auctioneer to personalize the resultant allocations (e.g. fair or robust allocations). Finally, we use the proposed framework to apply fairness mechanisms from a multi-attribute perspective in what we we call multi-dimensional fairness.

1.2 Contributions

The main contributions of this thesis are in the domain of artificial intelligence, being of special relevance for the development of multi-attribute resource allocation, mechanism design and multi-agent systems fields:

- Definition of an incentive compatible multi-attribute auction mechanism for resource allocation.
- Study of the requirements of an aggregation function in order that it can be used as evaluation function in a multi-attribute auction.
- Classifying the types of attributes that may take part in a multi-attribute auction.
- Definition of a framework for multi-attribute auctions in order to allow the utilization of different auction properties (e.g. incentive compatibility, fairness or robustness).
- Definition of a multi-attribute fairness mechanism based on the previously defined framework.
- Study and analysis of egalitarian approaches for the new fair auction mechanism.

1.3 Thesis Outline

- **Chapter 1: *Introduction*.** This chapter offers an overview of this thesis, its motivations, contributions, its methodology and the outline of the individual chapters.
- **Chapter 2: *Auctions for multi-attribute resource allocation*.** This chapter offers a review of the literature needed to understand this thesis. Within this chapter, previous existing auction mechanisms are described and analyzed, with criticism of their strengths and weaknesses.
- **Chapter 3: *PUMAA: Preserving Utility Multi-Attribute Auctions*.** A new multi-attribute auction mechanism based upon score functions is presented. We define the mechanism and analyze the requirements that an aggregation function must commit in order to be used in a multi-attribute auction. Moreover, the incentive compatibility of the mechanism is demonstrated.

-
- **Chapter 4: *A framework for Multi-Attribute Auction Customization.*** Within this chapter the attributes involved in a multi-attribute auction are analyzed and classified according to their roles and sources. This classification is used to extend the mechanism presented in the previous chapter in order that it can begin to accept other auction properties such as fairness. The extension of this mechanism is then demonstrated by defining a priority-based multi-attribute mechanism which can achieve fairer allocations.
 - **Chapter 5: *Multi-dimensional fairness for multi-attribute resource allocation.*** As the final methodology chapter, this chapter analyzes the requirements of fairness in multi-dimensional domains. We present a multi-dimensional fairness mechanism which we use to obtain egalitarian allocations using PUMAA.
 - **Chapter 6: *Experimentation and Results.*** In this chapter the methods presented in chapters 3 to 5 are tested in a simulated manufacturing environment, where agents in charge of different tasks have to perform them on-demand and accomplish production requirements.
 - **Chapter 7: *Conclusions.*** This chapter summarizes and discusses the research conducted in this thesis. Moreover, it suggests future research and improvements in systems which can be derived from this work.

AUCTIONS FOR MULTI-ATTRIBUTE RESOURCE ALLOCATION

This chapter provides background information regarding auctions and a brief overview of the state-of-the-art regarding auction protocols used for resource allocation. First of all, Section 2.1 introduces background information concerning auction and game theory, then Sections 2.2 and 2.3 describe the most relevant uni-attribute and multi-attribute auctions with regard to this thesis (Figure 2.1). Finally, Section 2.4 summarizes and discusses the current state-of-the-art.

2.1 Introduction

An auction is a protocol for buying and selling goods using a bidding system in which the winner bids obtain the auctioned goods [41] and where the price of goods is determined in the auction. In the manufacturing domain, auctions often follow a reverse schema: the auctioneer wishes to buy a service and the bidders are selling their working capacity at a given price. For the sake of clarity, we will specifically follow that approach along the thesis.

According to this approach (reverse auction), the main steps of an auction are as follows:

1. Auction set up or call for proposals (CFP): The auctioneer (a_0) sends a call for proposals to all of the participant bidders (a_1, \dots, a_n), in which the item(s) it to be bought (e.g. a task to be done) is described.
2. Bidding: agents acting as bidders return a bid to the auctioneer, with the price they are offering for the goods being auctioned, b_i . The price may or may not be equal to the

bidders ideal value for the item, depending on the bidder strategy or policy. The value that a bidder a_i gives to an item can be defined using a function $v_i(it)$. In an incentive compatible mechanism it is expected that bidders bid truthfully, so $v_i(it) = b_i$.

3. Winner determination problem (WDP): the auctioneer decides the winner(s) that maximizes its utility value [45]. The *utility* u_i is the measurement of the satisfaction received by the participants of an auction, either the bidders or the auctioneers [59]. Auctioneer's utility u_0 (Equation 2.1) is related to the payment and the item value: the auctioneer pays an amount p for an item it when the winner bid is b_i . Furthermore, the auctioneer has a value $v_0(it)$ for the item related to its interest in obtaining it.

$$u_0(it, b_i) = v_0(it) - p \quad (2.1)$$

Accordingly, the WDP can be formulated as:

$$\operatorname{argmax}_{i>0}(u_0(it, b_i)) \quad (2.2)$$

as the payment p_i is conditioned by b_i and it is constant, it can be simplified to:

$$\operatorname{argmin}_i(b_i) \quad (2.3)$$

4. Payment: the auctioneer pays the money to the winner(s). The utility of the winning bidder can be defined as its *profit*, so

$$u_i(it, b_i) = p_i - v_i(it) \quad (2.4)$$

where b_i is the bid submitted, $v_i(it)$ the value it gives to the item it and p_i the payment received. It can be observed that p_i is not necessarily equal to b_i since it will depend upon the payment mechanism established in the auction, as explained in the auction taxonomy below.

2.1.1 Basic game theory concepts

Auctions are a particular application of mechanism design [83]. Therefore, some basic game theory notions and definitions are required.

Mechanism design is a field of game theory that explores the outcomes in transactions among self-interested entities which have different goals, interests and private information. The goal of this field is to design games or mechanism maximizing the outcomes for the players involved

whilst inducing participants to share their private information in order to gain optimal outcomes. Therefore in the course of this thesis, when players are discussed, reference is made to the auction agents which act as auctioneers or bidders.

Throughout this thesis there are certain game theory concepts that will appear repeatedly:

- **Dominant Strategy:** In game theory one strategy is said to be dominant if for a given player that strategy is better than any other existing strategy, no matter which strategy the rest of players are using. Depending upon whether the strategy produces strictly higher outcomes or equal outcomes to other strategies the dominance is said to be *strict* or *weak*. In certain games or for certain players there might not be a dominant strategy in existence.
- **Nash Equilibrium:** A Nash equilibrium is produced when, in the course of a game, all the players follow a dominant strategy. In a Nash equilibrium, if a player changes its strategy it will see the quality of its outcomes reduced.
- **Incentive Compatibility:** A mechanism or a game is said to be incentive compatible, if for any participant, truth revelation is its dominant strategy. In other words, in a Nash Equilibrium all players reveal their true preferences.
- **Bayesian Nash Equilibrium:** When a player does not have a dominant strategy, the player behaves in a way which depends upon the strategy it thinks the other participants follow. A Bayesian Nash equilibrium is produced when all the players of a game follow the strategy that maximizes their outcomes, given their beliefs concerning other players strategies.
- **Bayesian Nash Incentive Compatibility:** This weaker notion of incentive compatibility is given when the best strategy for all the players is to reveal their preferences, assuming that the rest of players are also following this strategy.
- **Pareto Optimal:** An allocation is said to be Pareto optimal if there is no allocation which improves at least one player's outcome without worsening the rest. In multi-criteria decisions one configuration is said to be Pareto optimal, if there is no configuration that improves at least one attribute without worsening the rest.

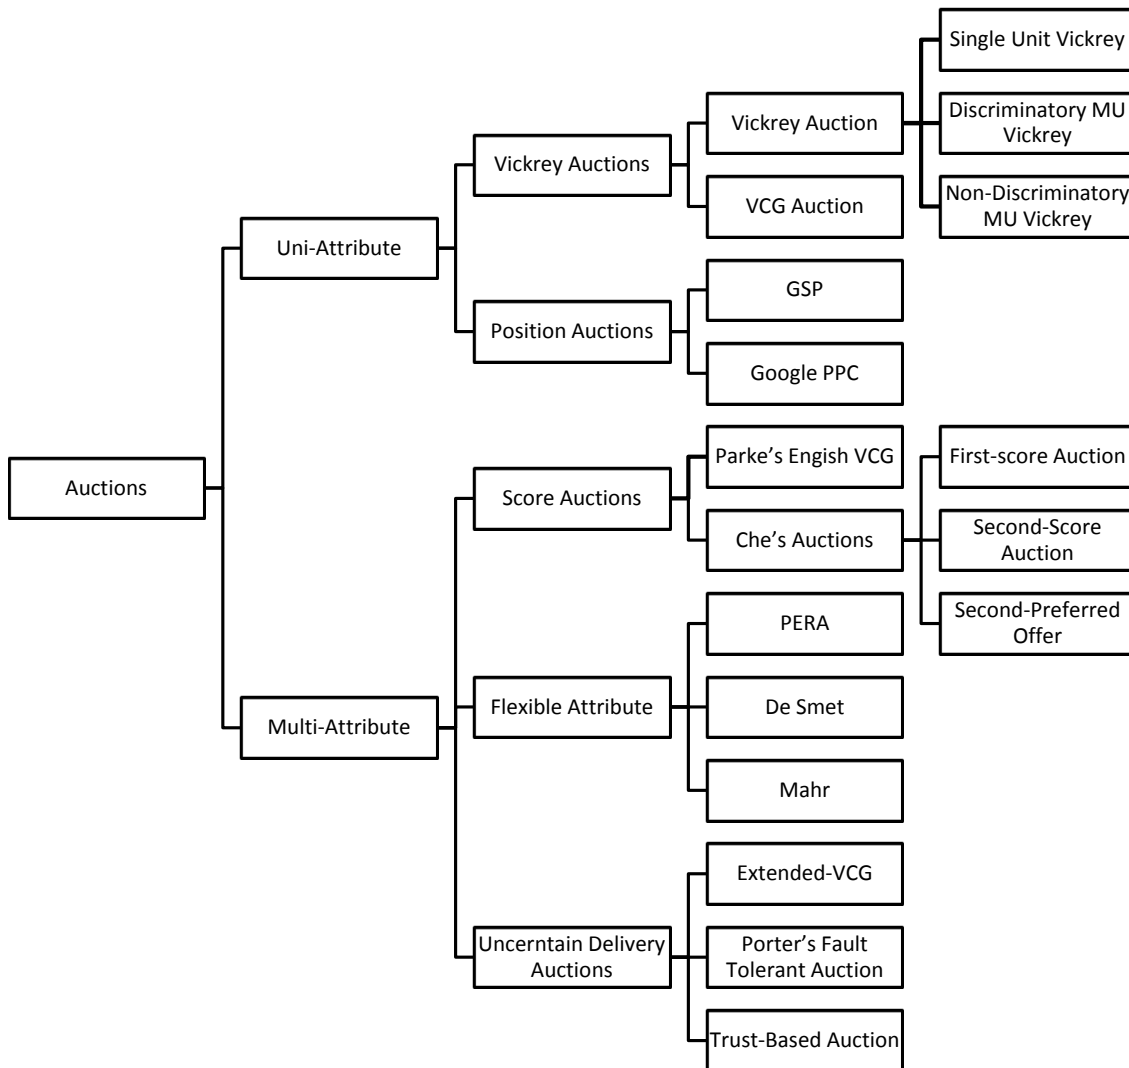


Figure 2.1: Auctions described in this chapter.




Figure 2.2: Auction classification according to bidding sides, sold resources, bid composition and the role of the participants.

2.1.2 Auction types

As Figure 2.2 presents, there are many auction types, which can be classified upon several criteria [62]. First, when a participant can play the role of auctioneer and bidder at the same time, this is double side auction. Conversely, in a one-side auction a participant can act only as the auctioneer or bidder. Second, the number of goods or services that are being sold in an auction determines if an auction is a single-item or a multi-item (combinatorial) auction. If the items auctioned are divisible (meaning the item can be split between two or more players) or indivisible (the auctioned item cannot be split) this changes the nature of the auction. Third, the number of attributes of the bid defines if an auction is a single-dimension (also known as uni-attribute) or a multi-attribute auction. In single-dimension auctions the bid is composed just by price whilst in multi-attribute auctions bidders provide information concerning other attributes beside the bid amount (e.g. time or quality). Fourth, if the winner of an auction pays the exact amount it offered in its bid, the auction is a first price auction; otherwise, if the price is conditioned by the non-winning bids it is a second price auction. Finally, according to the role of the participants, an auction can be classified as forward or reverse. As stated above, in forward auctions the auctioneer is selling goods or services and the bidders compete for them, in reverse auctions the auctioneer is the buyer and the bidders compete to sell their services.

In addition to that classification, auctions can be open-cry or sealed-bid [10]. In open-cry auctions, participants bid openly against each other, making their bids public and revealing their preferences to the rest of participants. In these mechanisms, agents privacy is sacrificed for the sake of transparency. A typical example of these kind of auctions are iterative auctions where the participants can modify their bids publicly depending on the other participants' bids. If in the course of the auction bidders are increasing their offers, this is an English auction, otherwise, in Dutch auctions bidders reduce their offers downward. Conversely, in sealed-bid

auctions, the privacy of the agents is preserved since all bidders simultaneously submit sealed bids to the auctioneer. In this way, no bidder knows how much the other participants have bid.

2.1.3 Auction properties

Given that auctions are a special case of mechanism design, auction mechanisms can be evaluated in terms of different properties typical of mechanism design:

- **Incentive compatibility:** A mechanism is said to be incentive compatible or strategy proof if all the participants obtain the highest utilities when revealing their real values, if it is assumed that all the participants are acting rationally [57] (in the Nash Equilibrium, truth revelation is the dominant strategy). If truthtelling is only weakly dominant (there are other strategies which are as good as truth revelation), the mechanism is said to be weakly incentive compatible. In auctions where (weak) incentive compatibility cannot be reached, it is desirable that the dominant strategy for a participant is to reveal its true values if the rest of participants are also doing so (Bayesian Nash Incentive Compatibility).
- **Efficiency:** An efficient task allocation is one in which the task to be allocated is assigned to the agent which values it most, and which has the highest capacity to perform the function. An efficient allocation in a game (like an auction) is also the one that maximizes the aggregated players' utilities under the game rules [42]. Some authors also consider an auction mechanism to be efficient if the winning bidder is not dominated (there exists no bidder in the game which can provide a better bid than the winning bid except the winning bidder itself) [8].
- **Buyer optimality:** In a buyer optimal mechanism the buyer pays the minimum possible amount for the item/allocation it desires. No second price auctions satisfies this property as the buyers may pay a higher price than the asked by the bidders. However, generally it is not possible for a buyer-optimal mechanism to be efficient and incentive compatible [49].
- **Individual-rationality:** A mechanism is individual-rational if the participants obtain equal or higher utilities by entering the game than by not participating. This means that rational agents do not obtain negative utility by entering the game. In auctions, a mechanism is individual-rational if the bidders do not obtain negative utilities by the

simple fact of providing a bid and if the auctioneer only sells/buys the auctioned items if it obtains a positive utility for it.

- **Budget-balance:** In a budget-balanced mechanism the transfer between players is equal to 0, in other words, the sum of the money the buyers pay must be the same that the sellers receive, without the need of third-party subventions. An auction is budget-balanced if, when considering the auctioneer as a participant, the sum of the outcome transfers of all the agents is equal to 0.

Moreover, auctions can be analysed from the point of view of the allocation in terms of:

- **Social Welfare:** It refers to the welfare of the participants involved in the allocation, this can be studied from different points of view: the utilitarian view favors the maximization of the welfare of all the individuals in the society, without taking into account the differences that are produced between the different participants. The egalitarian view approaches social welfare considering the sum of the individual's welfare but also taking into account the differences that appear between individuals and the envy which can appear amongst individuals [27]. Social welfare is considered fair if it maximizes the utilities of the participants whilst minimizing the disequilibriums and envies between them.
- **Robustness:** According to Holland and O'Sullivan [37] an auction with a robust solution is one that can withstand a fault (a bid withdrawal or a performance drop) by making changes in order to maintain the auctioneer's utility or, at least, to minimize its utility loss. Other definitions of robustness are possible, for instance Bofill et al. [12] and Zhang et al. [89] describe stronger notions of robustness in which they provide alternative allocation for bid withdrawals. Along the thesis we use Holland's definition which is focused on preserving the auctioneer's utility.
- **Reliability:** It defines the confidence an auctioneer can have that the allocation resulting from an auction is going to be accomplished. Thus an auction which favors trustworthy bidders will be more reliable than one which treats everyone equally.

2.2 Uni-attribute Auctions

Uni-attribute auctions are those in which the offers provided by bidders are strictly defined by one attribute (generally economic cost). This section offers a brief description of some of the

most common auction mechanisms (Figure 2.1).

2.2.1 Vickrey auctions

Vickrey auctions are probably the most popular version of second price-auctions. This style was originally designed as a single item auction, but was later generalized for combinatorial auctions.

Vickrey auction

In the course of a Vickrey Auction bidders submit their offers as a secret (or sealed) bid in order to obtain a single indivisible item. The winner of the auction is the participant that submitted the highest offer; however, the price paid is the that of the second highest-bid.

The Vickrey auction is an efficient mechanism as the winning bidder is the one providing the best value for the auctioneer. Moreover, it is also incentive compatible as bidding truthfully is the dominant strategy [85]: overbidding (bidding a higher price than the bidder's true value) may result in negative utilities, whilst underbidding (bidding a lower price than the bidder's true value) reduces the chances of winning the auction.

Example 2.1. *An auctioneer pretends to sell a used car, it calls for an auction and receives three different bids from agents a_1 , a_2 and a_3 : $\langle a_1 = 1,200, a_2 = 1,300, a_3 = 1,350 \rangle$.*

The car will be sold to a_3 which has offered the highest bid (1,350€) however he will pay 1,300€.

Multi-unit Vickrey auction

Vickrey auction can also be used to sell more than one item of the same type (k units of the same item) [11], e.g. tickets for a music concert. In this situation, the auction has k winners and it can be considered as discriminatory or non-discriminatory according to the way the payment is executed. If all winning bidders pay the price bid by the first not winning bidder ($k + 1$), the auction is considered non-discriminatory as all the winning bidders pay the same price. If the bidder offering the i highest bid pays the price offered by the $i + 1$ highest bid, it is said to be discriminatory. It is important to note that when the multi-unit sale is performed under a discriminatory schema, the mechanism loses its incentive compatibility. This is because all the winners except the final bidder would have paid lower amounts if they

had offered lower bids. Therefore, for some bidders truthful bidding is not the best strategy as they would have obtained higher utilities by underbidding rather than by revealing their true values. This situation is produced because despite all the units auctioned having the same value, they are sold at different prices.

The Vickrey auction can be generalized in order to sell multiple items of different types within the same auction. This combinatorial generalization is known as the Generalized Vickrey Auction or the Vickrey-Clarke-Grove (VCG) auction [17, 33].

VCG auction mechanism

The Vickrey-Clarke-Grove auction mechanism is a generalization of the Vickrey auction for multiple items, in which the winning auctioneers are determined by choosing the bid combination that maximizes the auctioneer's utility. In this case the auctioneer pays exactly its marginal contribution to the final auctioneer's utility [3]. The payment can be understood as the opportunity cost (this is the harm a winning bidder causes to the rest of the participants with its bid).

In this mechanism truthful bidding remains the dominant strategy as the strategies of overbidding and underbidding have the same results: increasing the chances of losing utility [6].

The payment rule for a winning bidder a_i in the VCG mechanism is defined as the sum of the prices the bidders would have paid if a_i had not participated in the auction, minus the sum of the payment that all the bidders except a_i would make when a_i participates in the auction.

Thus, if we define $B = [b_1..b_n]$ as the set of bids received during the auction, B^* as the set of auction winners when receiving B , B_{-i} as the set of received bidders minus the bid b_i delivered by agent a_i ($B_{-i} = B - [b_i]$) and B_{-i}^* as the set of auction winners when the auctioneer receives the set of bids B_{-i} we can define the payment rule as:

$$p_i = \sum_{b_j \in B_{-i}^*} b_j - \sum_{b_j \in B^* | j \neq i} b_j \quad (2.5)$$

Below we present a simple example illustrating how the VCG auction mechanism works and a second example showing that when only one item is auctioned VCG is analogous to the Vickrey auction.

Example 2.2. *An auctioneer pretends to sell two items: a computer and a flat-screen monitor. He summons an auction following the VCG mechanism. There are three buyers (a_1, a_2 and a_3) interested in buying these items: a_1 is interested in buying the screen for 100€, a_2 is interested in*

	Computer	Monitor
A_1	0	100
A_2	320	78
A_3	399	

Table 2.1: Table of bids corresponding to Example 2.2.

buying the computer for 320€ or the flat screen for 78€ (or both for 400€) whilst a_3 is interested in buying the whole lot for 399€, however he is not interested in the monitor or the computer separately. The received bids are represented in table 2.1.

Following a VCG auction the winners of the auction will be a_1 , who will get the flat screen monitor with a bid of 100€ , and a_2 who will get the computer with a bid of 320€ as this is the bid combination that maximizes the auctioneer's utility ($u_0 = 320 - v_0(\text{computer}) + 100 - v_0(\text{monitor})$). So, the payments they will incur are computed as follows:

- a_1 : The payment a_2 and a_3 must make when a_1 is participating in the auction is 320€ (paid by a_2), if a_1 had not participated in the auction the optimal allocation would have been a_3 , paying a_2 and a_3 a total of 399€. Thus a_1 's payment is 79€ ($399 - 320$).
- a_2 : a_1 and a_3 make a payment of 100€ (paid by a_1), when a_2 participates in the auction, whilst the optimal allocation removing a_2 allocates the whole set to a_3 for 399€. Thus a_2 will have to pay 299€ ($399 - 100$).
- a_3 : Since a_1 and a_2 payments are the same, whether a_3 participates in the auction or not (320 and 100), a_3 's payment is 0€ ($420 - 420$).

Example 2.3. The seller from the previous example realizes that he has forgotten to sell the keyboard he was using with the computer. Therefore, he decides to summon a new VCG auction selling the keyboard and receives the following three bids: $\langle a_1 = 10, a_2 = 12, a_4 = 8 \rangle$.

Following the VCG mechanism, the keyboard will be allocated to a_2 and he will make a payment of 10€ (the payment of bidders a_1 and a_4 are 0 when a_2 takes part of the auction and 10€ when he does not participate). If the auction follows a Vickrey schema, the keyboard would also have been allocated to the buyer a_2 for 10€ (the amount offered in the second highest bid, 10€ , provided by a_1).

The image shows a Google search interface with the query 'buy smartphone'. Below the search bar, it indicates 'About 335,000,000 results (0.25 seconds)'. The results are divided into two main sections: 'Sponsored/Ads' and 'Search Results'.

Sponsored/Ads:

- Slot 1:** 'Smartphone Deals' from phones4u.co.uk. It features a star rating, a link to the website, and promotional text: 'Get Your New Smartphone On An Amazing Contract Deal at Phones 4u!'. Below this, it lists 'Galaxy S4 Free From £33pm' and 'iPhone 5 Free From £37pm'. To the right, it says 'HTC One Red - Exclusive Upgrade Price Promise'.
- Slot 2:** 'New HTC One - Top Ratings - htc.com'. It includes a link to the HTC website, a star rating, and text: 'Stunning Camera, Fantastic Imaging Tools. Read our Customer Reviews! HTC has 172,496 followers on Google+'. Below this, it says 'HTC One Smartphone - HTC Sense BlinkFeed™ - User Ratings & Reviews'.
- Slot 3:** 'Top 5 Best Smartphone Deals - Compare and Save Money Today' from compareuk.net/Contract-Phones. It includes a link to the website and the text 'Everything Must Go!'.

Search Results:

Related searches: [buy smartphone unlocked](#) [buy smartphone handset only](#)

- [How to buy the best mobile phone - Mobile phone reviews - Phones ...](#)
- [www.which.co.uk/technology/.../how-to-buy-the-best-mobile-phone/](#)
- Mobile phones: How to **buy** the best mobile phone. How to choose your perfect mobile phone; **Smartphone** operating systems uncovered; Expert advice on the ...
- Choosing your perfect phone - Buying the best smartphone - Buying the best iPhone
- [Amazon.co.uk: SIM-free Mobile Phones & Smartphones: Electronics ...](#)
- [www.amazon.co.uk/Mobile-Phones/b?ie=UTF8&node=356496011](#)
- Results 1 - 24 of 7600 - Amazon's Picks-Find the Right **Smartphone** For You Eligible for FREE Super Saver Delivery and 1 more promotion; More **buying** ...
- [HTC Wildfire S Sim Free ...](#) - [Nokia N95 Sim Free Mobile ...](#) - [Nokia 2720 Fold](#)

Figure 2.3: Example of a Google sponsored search result.

2.2.2 Position auctions and the sponsored search market

In a position auction an auctioneer pretends to sell a given number of ranked slots or positions (for example the advertisements which appear in a web page) to a set of bidders. There are as many winners as slots auctioned and each winner is allocated according to the rank of its bid (the best bid obtains the first or the best slot whilst the i^{th} best bid obtains the i^{th} best slot).

Generalized second price auction

This problem is usually solved by using the generalized second price (GSP) mechanism for position auctions; an intuitive extension of the discriminatory Vickrey auction. Following the Vickrey philosophy, in GSP, each winner of the auction pays the amount it should have tendered to overcome the best bid after its own offer. Thus, the i^{th} best bid obtains the i^{th} position and pays the amount indicated in the $i + 1^{th}$ best bid.

This kind of auction is commonly used in web search engines to place advertisements alongside search results [41] (Figure 2.3). Given a set of keywords, the search engine desires to put

Bidder	Truthful bidding						Untruthful bidding					
	$v(i)$	α_i	b	Slot obtained	p_i	$u = \alpha(v - p_i)$	$v(i)$	α_i	b	Slot obtained	p_i	$u = \alpha(v - p_i)$
a	7	1	7	1	5	2	7	0.4	4.9	2	1	2.4
b	5	0.4	5	2	1	1.6	5	1	5	1	4.9	0.1
c	1	-	1	-	-	0	1	-	1	-	-	-

Table 2.2: Bids and their resulting utilities for Example 2.4

k advertisements before showing search results. Companies bid in order to appear in those slots. As the first slot has a higher probability α of being clicked than those which follow it, companies are interested in obtaining the first positions. Thus, if α_i is the probability of an advert placed in slot i being clicked, the utility of a bidder agent a which obtains the slot i can be defined as follows:

$$u_a(i) = \alpha_i(v_a(i) - p_i) \quad (2.6)$$

where $v_a(i)$ is the value the bidder gives to appearing in the search results at the i^{th} slot.

Despite this mechanism, which maximizes social welfare and is proven to be envy free (no bidder envies the position of the rest of competitors) [25], it is not incentive compatible as under certain circumstances truthful bidding is not the dominant strategy.

Example 2.4. Consider two advertisement slots with $\alpha_1 = 1$ and $\alpha_2 = 0.4$ and three bidders whose valuations are $\langle v_1 = 7, v_2 = 5, v_3 = 1 \rangle$.

In this example, truthful bidding for the participants is 7, 5 and 1. However, as can be seen in Table 2.2 this bid configuration is not in a Nash equilibrium as the first bidder would have obtained a higher expected utility by bidding 4.9€ than by bidding truthfully.

There is a variation of positions auctions for Internet advertising in which, rather than paying a determined amount for the slot obtained, the bidder pays an amount for each time the advertisement is clicked. This variation is known as pay-per-click position auction [84].

Pay-per-click position auctions: Google's approach

In pay-per-click (PPC) position auctions bidders offer the amount they aim to pay each time their advert is clicked. In this way, the bidder who obtains the first slot will be the one which offers the highest PPC and the i^{th} slot will go to the i^{th} highest PPC bid. However, this approach does not grant that the auctioneer's utility will be maximized.

Example 2.5. The web search engine "WeSearch" offers two ad slots during one week for the search keywords "Buy smartphone" with a probability of 0.5 and 0.25 of being clicked. There are three

	Slot	PPC	Quality q	PPC·q	Payment per click
myphone	1	0.050	0.600	0.030	0.017
unknownCo	2	0.100	0.100	0.010	0.020
unkownInc	-	0.020	0.100	0.002	-

Table 2.3: Resulting allocation Example 2.5 following the Google's pay per click position auction approach

websites interested in occupying those slots: two unknown electronic shops ("unknownCo" and "unknownInc") which offer 0.1 and 0.02€ per click respectively, and the famous and renowned cell phone manufacturer "myphone" which offers 0.05€ per click.

According to GSP position auctions, the first slot will be occupied by *unknownCo* and the second by *myphone* and they will pay 0.05 and 0.02 € per click (corresponding to the second and third best bids). Once the week has finished, "WeSearch" requires each company to pay for their ads. On the one hand *unknownCo* has received 1,000 clicks, despite being in the second slot, "myphone" received 4,500 clicks due to its fame. Thus, it receives a total revenue of 140€ (50 from *unknownCo* and 90). If we consider that being placed in the second slot reduces the probability of being clicked to half, is it reasonable to assume that *myphone* would have obtained many more clicks being in the first spot. In this scenario the auctioneer revenue would have been much higher (given that in this case α_1 is two times α_2 , we can estimate that if they had exchanged their slots *myphone* would have obtained approximately 9,000 clicks whilst *unknownCo* would have obtained 500 clicks).

To avoid that revenue loss, Google and Yahoo propose to add an attribute describing the likelihood of an advertisement being clicked given the quality of its design and the destination web page [84, 25]. This attribute is incorporated into the auction by the auctioneer itself and it is aggregated with the PPC offered by the bidders. The payment is then computed following a second price, the PPC of the i^{th} winner of the auction will correspond to the PPC it should have offered to equal the $i + 1$ best bid.

Example 2.6. If "WeSearch" considers that the quality for "myphone", "unknownCo" and "unknownInc" are $q=0.6, 0.1$ and 0.1 respectively then the resulting allocation is the one presented in Table 2.3. Thus, given the click probability for each slot and the previous estimations, the auctioneer revenue would be approximately 163€ ($0.017 * 9,000 + 0.020 * 500$).

It is important to note that the inclusion of this quality attribute does not affect the bidding strategy of bidders. The only influence of quality over bidders is that it encourages them to

improve their ads and their web pages in order to make them more attractive to the auctioneer so they can obtain a higher quality valuation.

Other authors have studied and designed mechanisms similar to this one. In [2] this mechanism is studied with the assumption that bidders may have information regarding the rest of participants' preferences and how the users of the web search can influence its results. The equilibrium present in these kinds of auctions is analyzed in [25]. Due to the difficulty of establishing a stable equilibrium, the authors transformed the GSP to an English auction in order to obtain analogous results and to analyze the possible bidding strategies which may result.

2.3 Multi-Attribute Auctions

In multi-attribute auctions bidders provide bids with an arbitrary number of attributes. These attributes usually include the economic cost and are complemented with different kinds of information, which may be physical properties of the item which is being auctioned (e.g. qualities, colors, etc.) or conceptual information (e.g. the probability of succeeding in performing a given task). In this section we briefly review some of the most popular multi-attribute auction mechanisms.

2.3.1 Scoring rules-based auctions

One of the main issues in multi-attribute auctions is defining how attributes should be aggregated and compared during the course of the WDP and the payment. In this section we describe some of the auctions which use an aggregation function known as *Score* which gives a real numeric value to the bid (containing the economic cost and the attribute bundle). This score is used to determine which is the best auction style and also, in certain mechanisms is used to compute the payment.

Che's First-score and Second-score auctions

One of the first authors to consider that an auction could be handled using more than one attribute was Che [15]. In his work Che studies auctions for government procurement where firms bid on price and other attributes. He aggregated these under a uni-dimensional quality attribute q , resulting in a bidimensional multi-attribute auction schema. To deal with this

dimensionality, Che introduces the usage of a scoring rule S which rates the received bids and which is used to determine the winner of the auction.

According to the author, the winner of the auction is determined by the scoring rule, which is used to determine the winner of the auction and must reflect the true preferences of the buyer which are defined by its utility function:

$$u_0(p_i, q_i) = v_0(q_i) - p_i \quad (2.7)$$

where p_i is the price the buyer pays for the auctioned item and $v_0(q_i)$ a function which values the quality of the item. Thus, if c_i corresponds to the economic cost which bidder a_i is asking for the auctioned item, and if $S(c_i, q_i)$ denotes a scoring rule for an offer (c_i, q_i) , the scoring rule can be defined, for example, as $S_0(x, y) = u_0(x, y)$ and the winner of the auction is the bid which maximizes S :

$$\text{argmax}_i S(c_i, q_i) \quad (2.8)$$

Using Equation 2.8 to solve the winner determination problem, Che proposes three different auction models which differ concerning the payment:

- **First-score auction:** In this auction type the winner receives/pays the amount indicated in the offered bid.
- **Second-score auction:** This auction, follows a Vickrey philosophy which forces the winner of the auction to match the score of the second best bid of the auction. This means that it can provide any pair (c_w, q_w) which provide the same score as the second best bid.
- **Second-preferred-offer:** is a particular case of the previous auction. Under this schema the winner is forced to provide exactly the same combination c_2, q_2 which is present in the second best bid.

Despite the Second-score auction being incentive compatible, taking into account that the best strategy for bidders is to provide their real values for the quality and price attributes, this situation is only accomplished if it is assumed that all the bidders follow the same behavior. The second-preferred-offer is not an incentive compatible mechanism despite it achieves Bayesian Nash Incentive Compatibility (the best strategy for an agent is to tell the truth if the rest of the participants are also revealing their preferences). This auction style can only be applied if it is assumed that all the bidders can provide the same range of qualities (assuming that if the second score has a higher quality than the first score, the agent winning the auction will be able to provide services of the same quality).

	Time (days)	Price (€)	Score
b_1	800	10,000,000	6,800,000
b_2	950	7,000,000	3,800,000
b_3	600	20,000,000	4,800,000

Table 2.4: Bids and scores for Example 2.7

The application of second-score auctions in the task allocation domain (by replacing quality for delivery times t) presents certain problems. The seller has complete freedom to provide any pair t, p matching the second best score making it impossible for the buyer to know when a task will be finished. This complicates the scheduling of future tasks due to uncertainty. On the other hand, despite the second-preferred-offer solves the uncertain delivery task issue, it is not an incentive compatible mechanism. Furthermore, it can not be assumed that two different agents will be able to provide the same delivery times since they can have very different configurations (different types of resources, different capacities, etc.).

Example 2.7. *The Government of Spain pretends to build a railway from Barcelona to Madrid. As government elections are in 3 years (1095 days), the government values the railway construction as 5 million of euros and, it wants the railway to be finished as soon as possible. It states that each day until the inauguration is worth 40,000€. However, due to the economic crisis, they want to spend as little as possible. With this in mind, the government summons an auction to decide who will develop the railway: the auction takes into account the cost of the works and how many days it will last. Given that the government utility function is defined as $U(p_i, t) = (5,000,000 + (1095 - t) \cdot 40,000) - p_i$, so it is its scoring function. The auction receives three bids from three different construction companies which bidded truthfully: $\langle b_1 = (800 \text{ days}, 10,000,000 \text{€}), b_2 = (950, 7,000,000), b_3 = (600, 20,000,000) \rangle$.*

Following the auction mechanisms proposed by Che the bid with the best score would be the first bidder (Table 2.4).

- If using a first-score auction, the winner would receive 10 million euros and would have to finish the works within 800 days. However, the winner would have obtained a higher utility if he had bidded a price of 11.9 millions because he would have won the auction and earned more money (the mechanism is not incentive compatible).
- In a second-score auction, the winner could finish the works within as many days as he desired and it would receive a payment corresponding to the score of the second best bid. For example:

- For finishing in 600 days, it will receive 20,000,000€ ($payment = -4,800,000 + 5,000,000 + (1,095 - 600) \cdot 40,000$)
- For finishing in 800 days, it will receive 12,000,000€ ($payment = -4,800,000 + 5,000,000 + (1,095 - 800) \cdot 40,000$)
- For finishing in t' days, it will receive p'_i ($p'_i = -4,800,000 + 5,000,000 + (1,095 - t') \cdot 40,000$)

In this case, the bidder obtains the highest utility by telling the truth.

- In a second-preferred-option auction, the winner has to finish the railway in 600 days and will earn 20 million. In this case it is assumed that the winner has the same capabilities as the second best bid. However, in reality, this assumption is too strong as there is no way to prove that the first bidder has those capabilities.

Che's mechanism for a 2-dimension multi-attribute auction is extended to an arbitrary number of auctions by **David et al.**¹ [21, 22]. In their work they also present an iterative auction model for multi-attribute auctions in which the agents could change their bids after every iteration as is done in English auctions.

Parkes' Modified VCG Auction

It is well known that for resource allocation decision problems the VCG mechanism may not be budget balanced (meaning that the transfer between the buyer and the seller is not equal, requiring a subvention for a third entity in order to cover the required costs) [20]. For instance, consider two neighbors who intend to buy an elevator which costs 10,000€ [30]. The neighbor from the first floor is willing to pay 7,000€ for it whilst the neighbor from the second floor is willing to pay 8,000€. The welfare value when constructing the elevator is 13,000€ (the sum of neighbor's valuations) whilst not constructing it has a welfare of 10,000€ (the money they save). Thus, following a VCG schema, the amount that first neighbor should contribute is 2,000 (10,000 - 8,000) whilst the second neighbor should contribute 3,000 (10,000 - 7,000), a total of 5,000€. However, this amount is not enough to build the elevator.

With this in mind, Parkes [60] considers that the VCG auctions (including the Vickrey specific case) are neither budget balance nor buyer-optimal. For instance, if we consider the second-score auction in Example 2.7, we can see that the Government pays 12 million euros when the bidder was prepared to do the work for 10 million. Another example of this imbalance is

¹Authors and methods marked in bold appear in Figure 2.6

produced in the following case: an auctioneer who pretends to sell a painting, which he values at 100€, could sell the painting in a Vickrey auction below this price despite having an offer higher than his own valuation. For example the best bid might be 105€ whilst the second best bid is 95€. With these bids the winner would pay 95€ for the painting and the auctioneer would obtain a utility of -5 despite receiving an offer which offered him a positive utility.

To deal with that problem Parkes presents a reverse iterative VCG auction modification which also uses a scoring rule but modifies the way the payment is computed. In the Parkes approach bidders have the chance to revise their offers iteratively once all the participants have placed their bids in a similar way to the English auction approach presented in [21, 22] by David et al. As in the previous approach the auction is cleared by means of a scoring function \hat{S} , however, unlike Che's auction, this scoring function does not have to correspond to the utility of the auctioneer. Moreover, this scoring function can change at each iteration in order to tune the results according to the auctioneer's preferences. Once bidders stop modifying their bids, the auction is cleared selecting the bid with the highest score. The payment is then computed using the reported economic cost c_i offered by the best bid plus the difference between the score of the winning allocation and the score the winning allocation would have obtained if the winner bidder had not participated (the second highest bid in the Vickrey auction).

$$p_i = c_i + (\hat{S}(t) - \hat{S}(t/i)) \quad (2.9)$$

where $\hat{S}(t)$ is the valuation of the winning bid b_i and $\hat{S}(t/i)$ is the valuation of the bid which would have won if b_i had not participated in the auction.

Under this schema, if the auctioneer uses its utility as scoring function, the winning bidder will receive the same payment as in Che's second-score auction. However, if the auctioneer slightly modifies its scoring function, the winner will receive less money and the budget imbalance will be reduced. The downside of this mechanism is that it loses its incentive compatibility on the auctioneer's side as it may obtain higher utility providing scoring functions different to its utility function. This method is revised and extended in another piece of work by the same author [61].

2.3.2 Auctions with flexible attribute structures

The mechanisms described above require that all the bids are composed by the same set of attributes so they can be compared. This situation requires that the attributes can be converted to a numerical value so that they can be ordered. Thus, these mechanisms are only valid for domains where the auctioneers have linear or quasi-linear utility functions and preferences.

To deal with domains where the auctioneers have non-linear preferences, some authors propose using a winner determination problems based upon partial relation-based preferences.

Preference-based English Reverse Auctions

Bellosta et al. [7, 8] propose adapting English auctions in order to allow these kind of non-linear preferences in what they call preference-based English reverse auction (PERA). In this scenario, a single auctioneer wants to buy a single unit of a given item. PERA modifies the classical reverse English auction algorithm [14] in two ways:

- It changes the numerical bid comparison (based only on price in uni-attribute auctions and score functions in multi-attribute auctions) to the buyers preference relation. In this preference relation, given two bids b_1 and b_2 which do not necessarily have the same structure, the buyer or auctioneer can prefer a bid to another (\succ) or have no preference among the bids (\approx). It is important to state that the auctioneer preference relation can be ordered ($b_1 \succ b_2, b_2 \succ b_3, b_1 \succ b_3$) but also partially ordered ($b_1 \approx b_2, b_2 \succ b_3, b_1 \approx b_3$).
- At each iteration, instead of accepting bids with a lower price (or score) than the provisional winner, the auctioneer only accepts bids which are preferred to the provisional winner's bid (new bid \succ provisional winner bid).

[8] also describes how other auction mechanisms can be modeled under the PERA schema. With respect to multi-attribute auctions, it describes how a score-based English auction mechanism can be converted to PERA as the score rule results can be used to model the preference relation among bids.

With this mechanism Bellosta et al. allow the comparison of bids which do not necessarily have the same structure whilst assuring efficiency (the winning bid is non-dominated by any other bid according to the buyer's non-linear preferences). The authors do not discuss the strategy proofness of the mechanism neither when using partial ordered preferences nor when using ordered relations. Moreover, PERA shares deficiencies with English auctions when used for long contract procurements. As these procurements involve a paused negotiation, in dynamic and fast domains where resource allocation must be fast and almost automatic, the iterative negotiation and the expression of the relations would greatly slow the process.

Similar approaches and cases of study can be found in the study of **De Smet** [75], where the author studied bids which do not share the same attributes in English auctions. This is

also studied by **Mahr and de Weerd** [50],], who propose an adaptation of Che's second-preferred-offer auction based on order of preferences, for cases where the auctioneer's utility is unknown; Mahr's approach is generalized in [34]. In these articles, as well as in PERA, the authors suggest the use of a non-numeric preference relation system for determining who the winners of the auctions are.

2.3.3 Auction-based allocations where the task delivery is uncertain

Many of the reverse auction mechanisms used to allocate tasks to external providers are based upon the strong assumption that once an auction is finished and an allocation is defined, all the participants will be able to succeed in executing their tasks. However this is an unrealistic assumption, as in many domains agents can fail to perform their tasks. For instance, a courier company may fail to deliver a packet, reducing the utility of its service to zero or even less. Diverse studies have tackled this issue by means of the inclusion of elements defining the chances of an agent being able to succeed in its task, such as the probability of success (POS), reputation or trust.

Due to this uncertainty regarding bidders being successful in performing their tasks, auctioneers cannot longer know the utility they will obtain from the resulting allocation prior to the end of the task. Thus, auctioneers have to determine the winner of the auction by predicting the probability of that agent succeeding in performing the task. Maximizing the auctioneer's expected utility.

The probability of success $\rho_j(T)$ is the probability of an agent a_j being able to perform satisfactorily given task T . Thus, we can define the expected utility $\bar{u}_i(T, \rho_j(T))$ of a task T as the combination of the utility obtained from a successful deployment of $u_i(T)$ and the probability of T being successfully performed:

$$\bar{u}_i(T, \rho_j(T)) = u_i(T) \cdot \rho_j(T) \quad (2.10)$$

VCG POS Extension

Is it reasonable to assume that each agent knows its own POS or, at least, it has an accurate estimation $\bar{\rho}$ of the POS in a given task. Taking this into account, an intuitive approach is to extend the VCG mechanism to include each agent POS in the bid [69]. Thus, if an auctioneer a_0 auctions a task T it will receive bids following this structure $b_i = (c_i, \bar{\rho}_i(T))$. The winners of the auction will be determined by the allocation Γ which maximizes the sum of expected utilities. The payment rule would be defined similarly to the VCG auction:

$$p_i(T) = \sum_{b_j \in B_{-i}} \bar{u}_j(\Gamma_{-i}) - \sum_{b_j \in B_{-i}} \bar{u}_j(\Gamma_i) \quad (2.11)$$

where B_{-i} represents the set of all bidders except the one (b_i) receiving the payment, Γ_i represents the winning allocation when agent b_i participates in the auction and Γ_{-i} when it does not.

Example 2.8. *The car manufacturer WeMakeCars wishes to send a brand new car to one of its authorized dealers. For that purpose WeMakeCars wants one of the many courier companies in town to transport the new car from the factory to the concessionaire. Each courier company has a different cost for performing the delivery (depending on the number of employees they have, the quality of their trucks, etc.) and a probability of succeeding in the delivery task (depending in factors such as the chosen route or the ability of their drivers) which are only known by the courier company itself. WeMakeCars obtains a utility of $u_0(T) = 100$ from delivering the car from the factory to the concessionaire. Under these conditions, WeMakeCars summons an auction and receives three bids $\langle b_1 = (30, 0.5), b_2 = (50, 0.9), b_3 = (70, 1) \rangle$ (see Table 2.5).*

In this example the winner of the auction is bidder 2 with an expected utility of $100 \cdot 0.9 - 50 = 40$ and the payment it will receive according to the extended VCG mechanism is $70 \cdot 1 - 0 = 70$. However, this VCG extension fails in terms of incentive compatibility because some agents may have obtained higher utilities by providing a false value for $\rho_i(T)$. For example, if bidder 1 had submitted a false $\rho_1(T)$ value of $\rho'_1(T) = 1$ the expected utility of its bid would have been $100 \cdot 1 - 30 = 70$ hence winning the auction and obtaining a payment of $50 \cdot 0.9 - 0 = 45$. Thus, the utility of bidder 1 would have been higher by providing a false ρ than by revealing its true value.

Porter et al.'s Fault-tolerant Mechanism

To face the problem of merging the POS with the VCG auction mechanism Porter et al. propose the incorporation of a set of modifications into the mechanism's structure [65]. First of all the

	c_i	$\rho_i(T)$	$\bar{u}_0(c_i, \rho_i(T))$
bidder 1	30	0.5	20
bidder 2	50	0.9	40
bidder 3	70	1	30

Table 2.5: Received bids by WeMakeCars and their expected utilities given that $u_0(T, c_i) = 100 - c_i$

payment step is moved to the end of the auction, thus, payment is not provided until the completion of tasks, in such a way that a Boolean variable k is added into the mechanism. This value k defines if a task has been successfully performed ($k = 1$) or if it has not ($k = 0$). This provides the auctioneer with the ability to pay or punish the bidder depending on whether the task has been accomplished.

The winner is determined in the same way as in the VCG trust extension, by choosing the allocation which maximizes the auctioneer's expected utility. However, this mechanism differs in the payment rule, if the bidder succeeds in the task performance ($k = 1$), following the VCG philosophy, the payment extracts the expected marginal contribution of the winner bidders by comparing the resulting allocation with the allocation it would have resulted if the winner bidders had not participated. If the bidder fails in the task ($k = 0$) the bidder is required to pay the expected utility the auctioneer has lost by allocating the task to the winning bidder.

$$p_i(T, \rho_i(T), k_i) = \begin{cases} \sum_{b_j \ni B_{-i}} \bar{u}_j(\Gamma_{-i}) - \sum_{j \ni B_{-i}} \bar{u}_j(\Gamma_i) & \text{if } k_i = 1 \\ -\sum_{b_j \ni B_{-i}} \bar{u}_j(\Gamma_i) & \text{if } k_i = 0 \end{cases} \quad (2.12)$$

which can be simplified as:

$$p_i(T, \rho_i(T), k_i) = \left(\sum_{j \ni B_{-i}} \bar{u}_j(\Gamma_{-i}) \right) \cdot k_i - \sum_{j \ni B_{-i}} \bar{u}_j(\Gamma_i) \quad (2.13)$$

If we use this mechanism to allocate the task proposed in example 2.8 when bidder 1 submits a false $\bar{\rho}(T) = 1$ then the bidder would have obtained a payment of 45 if $k = 1$ and a payment of -45 otherwise. If we take into account that bidder 1's utility is defined as the payment it receives (45 or -45) minus the economic cost of performing the task (30) we can see that the expected utility for bidder 1 is $\bar{u}_1 = 0.5 \cdot (45 - 30) + 0.5 \cdot (-45 - 30) = -30$. In the case where all the bidders have bid truthfully the expected utility for the winner (bidder 2) is $\bar{u}_2 = 0.9 \cdot (70 - 50) + 0.1(-70 - 50) = 6$.

This mechanism does not incentivize agents to lie regarding their POS, as reporting $\rho' > \rho$ increases the chance of being allocated despite not being the best choice. This increases the probability of receiving an expensive punishment in case of not being able complete the task. In other words, providing a false POS makes the bidder's expected utility negative. A complete treatment of the incentive compatibility proofs can be found in Porter et al.'s paper [65].

This mechanism is successful in encouraging agents to reveal both their real costs and POS, however, it has a few drawbacks. It does not take into account that different agents may have a different perception of what constitutes a successful task (see Example 2.9). It also requires

that agents pay a penalty or a fee when they are unsuccessful in performing a task whilst there is no guarantee that those agents will pay it. Moreover, despite being a multi-attribute auction, Porter's proposal only support two attributes (economic cost and POS).

Example 2.9. *Following Porter's mechanism, the manufacturer WeMakeCars assigned the task of transporting a new car from the factory to a concessionaire to TheCourierCompany. Despite TheCourierCompany believing that the task was accomplished as the car was successfully delivered into the concessionaire, WeMakeCars stated that the task as failed as the car was scratched and dented. Thus, WeMakeCars can not completely trust TheCourierCompany's POS because they have different metrics for evaluating the success of a task.*

Ramchurn et al.'s Trust-Based Mechanism

To avoid depending exclusively on the self-observed probability of success of each agent, Ramchurn et al. [69] propose a method where several agents may have estimations of other agents' POS (based on information obtained through past experiences or other sources) in a reputation model which they call *trust*. Estimations based on observations are typically affected by noise, however, auctioneers may obtain more accurate information concerning the quality of a certain agent's POS by gathering the observations obtained by the rest of agents and weighting the agents' estimations based on the similarity of their preferences.

For this purpose, agents record (from their own point of view) how well the rest of agents perform the tasks in their charge, in this way, each agent provides an opinion value called trust concerning how trustworthy the rest of agents are. In this way, when an auctioneer summons an auction, instead of taking into account the estimations each bidder has of its own POS, it asks to the rest of agents for their trusts and calculates the expected utility for each bid using value of the economic bid and the received trust values.

This mechanism achieves incentive compatibility regarding the economic bids using a VCG schema. In order to encourage agents to provide real values for their trust, the auctioneer rewards the agents which have offered useful opinions with a small amount of the payment. As well as Porter's auction, this auction only supports two preestablished attributes (economic cost and trust).

2.3.4 Attributes as constraints

Certain mechanisms which involve attributes other than the economic cost opt for a more simple solution using the attributes as a constraint.

In these cases the auctioneer fixes a range of values for the extra attributes and only accepts bids which are within the delimited values, if a bid does not fit the marked constraints it is excluded from the auction. When all the bids are filtered according to the delimited constraints, the only attribute used to choose the winner of the auction and its payment is the economic cost. This kind of mechanism considerably simplifies the auction process, however, it does not differentiate the quality of the attributes involved in the auction. It is a black and white method in which a bid is taken into account if it commits to certain requirements or it is ignored if the requirements are not fulfilled. For example, under a task allocation system which considers delivery time and cost, an auctioneer could fix a delivery time window between 20 and 40 minutes. All the bids within this time range will be considered whilst the bids with a delivery time lower than 20 or higher than 40 would be ignored. The winner then would be the bid with the best economic cost, regardless of whether the delivery time is 21 or 40 minutes.

This approach has been described by authors studying multi-attribute auctions, whose major concern was another aspect of the auction. For example double auctions with high computation complexity where bidders can re-auction the tasks they bid for. This is the approach followed by **Zhao et al.** [90], they present a reverse double auction mechanism to allocate time-dependent tasks. In their approach agents act as sellers and buyers at the same time. Each auctioneer asks for a set of tasks to be allocated within a time window and the bidders compete to perform those tasks, however, bidders can re-auction the tasks if they consider that this will generate them more profits. Many auctions are cleared simultaneously by a central agent using a graph-based algorithm which is focused exclusively on the price. Only the bids which commit the time requirements are taken into account.

Similarly, **MAGNET** [19] uses time constraints to determine workflow schedules with temporal and precedence constraints. Customers act as auctioneers and summon auctions in order to find supplier agents to perform specific tasks. On the other hand, supplier agents (bidders) submit bids specifying prices for combinations of tasks together with time windows and the duration of those tasks. As customers have temporal and precedence constraints, they must pick a combination of bids which satisfies their constraints. In this approach, using an heuristic search, time is used to determine which bids are feasible (respect the time constraints), but the winner is computed by minimizing the amount paid (the winners are the best bid combinations amongst the set of feasible bids).

Despite these kind of auctions deal with multi-attribute bids, their WDP does not face a multi-criteria decision problem like the multi-attribute auctions previously described. In this case, the core of the WDP is focused upon the the improvement of the economic aspect of the

	Uni-attribute auction	Multi-attribute auction
Uni-criteria WDP	Vickrey Auction Multi-unit Vickrey Auctions VCG Auction Generalized Second Price English Auctions Dutch Auctions	Zhao's Double Auction MAGNET
Multi-criteria WDP	Google PPC Auction	Che's Auctions Parkes Modified VCG David's English Auction PERA De Smet auction Mahr Auction VCG-POS Porter's Fault Tolerant Auc. Ramchurn's Trust Auc.

Figure 2.4: Winner Determination Problem dimensionality of the auctions described above.

bid as varying the value of the attributes does not modify the utility of the auctioneer whenever the attributes are within the delimited boundaries.

2.4 Summary

This section has presents a brief summary of the mechanisms described in this chapter. First we analyze their dimensionality and we subsequently describe the properties of each auction typology.

2.4.1 Dimensionality

Uni-attribute auctions base the allocation of goods in the negotiation of just one attribute, usually the economic price. Multi-attribute auctions take this concept one step further by taking into account multiple elements during the decisions taken during the different steps of an auction.

Previous works in the literature considers that multi-attribute auctions are mechanisms

which automate the negotiation at a multiple attribute level, meaning that the process of determining the winning allocation of the goods takes into account more than one attribute [9], e.g. price and quality. In other words, received bids are n-dimensional. This implies that the auction process, from the call for proposals to the payment, is a multi-criteria problem for both the auctioneers and the bidders.

There are, however, a subset of uni-attribute auctions which share some characteristics with multi-attribute auctions (e.g. Google's position auctions). In these auctions the auctioneer calls for a single-attribute auction, the bidders submit single attribute bids but then, the auctioneers incorporate one or more extra attributes into the process of determining the auction winners. This provides an improvement of certain allocation properties (e.g. fairness, robustness or quality). Despite not being multi-attribute auctions, these kinds of auctions deal with a multi-criteria WDP.

There are also multi-attribute auctions which, despite dealing with n-dimensional bids, only take into account a single attribute during WDP. Sharing the WDP schema with most uni-attribute auctions, this kind of auctions first apply a filtering to determine which bids are feasible according to certain constraints. Then the auctioneer determines the winner of the auction following only one criterion (usually the economic value).

Figure 2.4 classifies the auction styles described in the previous sections according to the dimensionality of their WDP. In this thesis we are specifically interested in multi-attribute auctions, however, the mechanism presented in Chapter 4 also allows the modeling of uni-attribute mechanisms with a multi-criteria WDP.

2.4.2 Auction properties

Figures 2.5 and 2.6 summarize all the auctions described in the previous sections. In this table we can observe some remarkable aspects which will shape the nature of the work presented in this thesis.

One of the first things we can observe is that buyer-optimality and incentive compatibility are two properties which are usually confronted, except in English-based auctions (Parkes' English Auction, PERA, De Smet auction and Mahr auction). Moreover, analysis shows that incentive compatibility can only be reached by mechanisms following a second-price (or Vickrey) philosophy or an English-auction structure (which theoretically, obtain allocations and payments equivalent to those obtained with second-price auctions [53]). Another interesting aspect is that almost any multi-attribute auction can be useful for negotiating long-term procurement




Figure 2.5: PUMAA shown with respect to the auctions described in this chapter.

problems, however, only sealed bid auctions seem to be suitable for dynamic task allocation. This is because open-cry protocols may be too slow and heavy to allow fast negotiations. One can observe that in task allocation auctions it is generally assumed that the task performer will fulfill the auction agreement and that it will succeed. Auctions which do not make this assumption are described in Section 2.3.3. In this case, auctions adopt a black and white position in which a task can be a success or a failure, without considering any intermediates.

In this thesis we present PUMAA auctions and the FMAAC framework (which uses PUMAA as its starting point for customizing multi-attribute auctions).

The protocols presented in this manuscript follow a Vickrey approach and are based upon the dynamic's of Che's second-score auctions. However, unlike Che's proposal, we do not assume that a bidder can offer exactly the same attribute configuration as the second best bid. In this sense, we fix the set of attributes provided in the winning bid and compute the payment using the second best bid. PUMAA and FMAAC are designed to be applied in highly dynamic contexts where task allocation must be performed automatically and rapidly. This requirement makes us discard Parke's approach which is more suitable for long-term contract negotiation. In the same way, this discards the approaches presented in Section 2.3.2; thus our approaches are designed for domains where all the agents bid using the same categories of attributes.

Regarding uncertainty, on the one hand, both PUMAA and FMAAC present a similar structure to Porter's auction. The payment is performed after tasks are complete and it is conditioned by the satisfaction of the task conditions. Moreover, both approaches can be considered robust in case of the failure of the bidder to execute the task, as the auctioneers are compensated by means of a payment modification. In this way, the auctioneer can react to the failure and fix the task development problem in the next auction. Conversely to PUMAA and FMAAC, Porter's auction requires the bidder to define its probability of success. This probability will condition the winner of the auction and the payment it will receive. PUMAA does not use this probability and computes the payment directly depending on the success of the task. Moreover, our approach can be used using several attributes besides the price, whilst Porter's can only deal with the economic cost and the probability of success.

Ramchurn's auction offers robustness and reliability based upon agent's reputation. However, the structure of the auction limits its dimensionality to the economic value and the agent's reputation. Despite PUMAA might incentivize auction participants to provide accurate attribute estimations, it can not be considered to be reliable as it treats all the agents in the same way (without considering if they have been reliable in past). However, PUMAA may be customized using FMAAC in order to introduce trust and to favor agents which, in past

auctions, displayed good performances.

Finally, it can be said that PUMAA and auctions which treat attributes as constraints treat a different problem. This kind of auction, despite dealing with more than one attribute, pretends to maximize utility only regarding the economic cost whilst PUMAA deals with a multi-criteria maximization. Moreover, the dimensionality problems each auction faces have different origins: whilst dimensionality in constraint auctions is due to dealing with combinatorial auctions, the dimensionality of PUMAA and FMAAC is caused by the multiplicity of attributes to be maximized.

PUMAA: PRESERVING UTILITY MULTI-ATTRIBUTE AUCTIONS

In this chapter we present a new multi-attribute auction mechanism specially designed for allocating tasks or resources under demand in a workflow domain: PUMAA (Preserving Utility Multi-Attribute Auctions).

During the allocation process, we are interested in qualitative attributes of resources in addition to price; for instance the delivery time of the task, its quality, or the amount of CO₂ it will produce. For the design of the new mechanism we propose, we have taken ideas from position auctions [84] (where the auctioneer adds quality to the price offered by bidders) and from Che's second-preferred-offer auction [15] (where the attributes are offered by the bidders). Despite position auctions being a good starting point for this purpose, Google deals with attributes provided by the auctioneer, not by the bidder. In consequence, it is assumed that the attribute values are always reliable. Similarly in second-preferred-offer auctions it is assumed that the attributes provided by the winner will be the ones provided in the second best bid, and that the winner will not provide a different set of attributes. In our problem, where attributes are provided by agents that may fail to commit their agreements (e.g. delivery of a task later than agreed), this assumption cannot be done. For instance, bidders could lie when providing the attribute values in order to increase their utility or they could suffer unexpected mishaps that could compromise the completion of a task under the agreed terms. Thus, we need to improve the second-preferred offer, whilst assuring truthful provision of attribute values.

Furthermore, besides preventing bidders from intentionally providing tasks under poorer conditions than those agreed [40], we want to reduce the harm that an auctioneer may suffer

if a task is not properly performed (e.g. a task is delivered later than agreed or with lower but acceptable qualities). In other words, we want to preserve the utility of the auctioneer in the case where a bidder fails to carry out its bid. This problem is addressed by means of a conditional payment in which the bidder will receive an amount depending on the conditions of the delivered task.

Considering all these aspects, we present a new Vickrey-based reverse multi-attribute mechanism for allocating tasks to resource providers considering that the resource providers may fail to successfully perform the allocated tasks: Preserving Utility Multi-Attribute Auctions (PUMAA). The main characteristics of PUMAA are the minimization of auctioneers' utility loss when bidders fail to accomplish their tasks whilst ensuring truthful bidding is the dominant strategy for bidders.

3.1 Assumptions and Limitations

Our proposal involves the use of an auction mechanism for deciding which tasks to allocate to the available resource providers in a business process. Due to the uncertainty involved in business processes (decision points), the tasks that need to be executed are known not long before they are about to start. With this in mind, we interleave resource allocation with task execution, with a realization of the resource allocation on-demand. Furthermore, in workflows, underperformance of a task may produce undesired consequences (e.g. delivering a task later than expected may reduce the time left for performing the following task, thus, increasing the expense of executing the workflow on time). To minimize the effect of these unexpected and undesired situations, PUMAA offers a payment mechanism that minimizes the auctioneer's utility loss in the case that an outsourced task is not performed as agreed (e.g. reducing the payment to the bidder in order to have more budget for hiring a faster resource provider for the incoming task).

Other approaches are possible, for instance different sets of tasks could be allocated at the same time, however we prefer to interleave the allocation with production since this methodology does not cause overlap and pre-booking situations with resources, which could result in failures at run time. Moreover, interleaving production task allocation favors the flexibility of production, allowing manufacturers to follow production philosophies such as *Lean manufacturing* [73] or just-in-time (JIT) approaches like *Toyotism* [48]. Furthermore, linking the resource allocation to an under demand production allows the supply of more customizable items and a reduction of stocking levels.

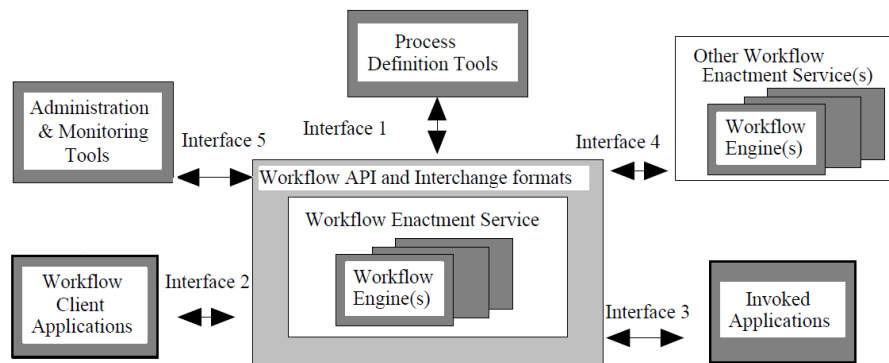


Figure 3.1: Workflow reference model according to [87].

The auction mechanism we propose is intended to be embedded in workflow management systems that take care of the operational issues of the business process (see Figure 3.1). The mechanism is especially appropriate for multi-agent workflow management systems [63] that control different workflows and providers by means of agents (see Figure 3.2). We assume that each business process is being handled by a business process agent (BP agent), whilst resource providers are represented as resource provider agents (RP agents). Thus, agents allow to represent all the participants of the system, whether they belong to third party companies or not [88].

When a business process is enacted, the corresponding BP agent monitors and manages its development. When it detects that a task needs to be externalized the BP agent auctions the task, assuming the role of the auctioneer. Resource providers which desire to perform the task, assume the role of the bidders. Once the task has been performed the workflow procedure continues until a new task needs to be outsourced, in which case a new auction can be called and so on. Therefore PUMAA is sequential, as the valuations of bidders can be influenced by past allocations and the decision to submit a bid to an auction or not will condition the following actions of the resource agents [39]. For example, an agent that participates in an auction may win it, thus, the agent compromises its occupancy during the time it takes to complete the task. This situation could prevent the winning agent from participating in the next auction, which could be more profitable for him (the agent has already committed its occupancy) as during this thesis we assume the use of non-preemptive tasks. Dealing with strategic issues regarding sequential auctions is outside of the scope of this thesis, and we assume that resource agents who are interested in the task auctioned will participate in the auction, however agents may learn from an auction to the next which strategy will improve their bids and profits (they can learn and modify their bidding strategy).




Figure 3.2: Multi-agent system schema, each business process is monitored by a BP agent while each resource is represented by a resource agent.

Along the thesis we assume that the variation of the attributes of a task conditions the utility of the auction participants and that the auction participants can tolerate such variability. For BP agents (auctioneers) this implies that receiving a task with quality attributes different than the expected will vary their utilities (for better or worse). On the bidders' side, we also make this assumption and we consider that providing a set of attributes or another will condition their economic true-values and, consequently, their utilities. This assumption also entails that, in order to ensure a proper quality of service, an auctioneer must include all the aspects that condition its utility as attributes of the auction (for instance, delivery time or product quality).

3.2 Mechanism

Given the auction classification proposed in Chapter 2, this mechanism is a sealed bid, one-side, single-item, multi-attribute, second-price reverse auction. The auction is one-sided given that auctions are solved one at a time and that agents can only play one role during the auction; single item since only one task is auctioned at every auction; multi-attribute because the auction winner and its payment value is decided by taking into account the economic cost but also a set of attributes; second price because the payment is based on the Vickrey's payment philosophy; reverse since the auctioneer is acting as the buyer (a BP agent that pays another agent for deploying a task) and bidders as sellers (they offer their capacity of performing tasks in exchange for money). It is a sealed bid auction as bidders can only provide one offer which can not be modified and that will remain secret to other participants.

The steps followed in PUMAA each time an agent needs to outsource a task are as follows:

1. **Call for proposals.** An agent needs a resource to deploy or externalize a task T^j , and it calls an auction for that purpose. The resource and task requirements are characterized by a set of numerical attributes requirements $AR = \langle ar^1, \dots, ar^n \rangle$ (e.g. earliest starting time t , quality q , etc.). Thus the agent calling the auction becomes the auctioneer a_0 , the task to be allocated becomes the auctioned task T_0^j and the task requirements defined by the auctioneer become AR_0 .
2. **Bidding.** The agents that can fulfill a_0 requirements can participate in the auction as bidders $a_i | i > 0$. Bidders will submit their bids $B_i = (b_i, AT_i)$ where b_i is the economic cost (or the price of the task) and $AT_i = \langle at_i^1, \dots, at_i^n \rangle$ the attribute qualifications.
3. **Winner determination.** The auctioneer a_0 evaluates the bids using an evaluation function $V_0(B_i)$ according to the auctioneers utility. As we are dealing with a reverse auction, the lower values of V_0 represent a higher utility for a_0 . a_0 cleans the market by ranking the bids from the lowest to the highest value. The bid with the lowest value is the winner of the auction.
4. **Payment.** When the winning agent a_1 completes the task, it receives a payment p_1 . The payment will depend on whenever the task has been delivered according to the attributes it bidded.

These steps are detailed in depth below.

3.2.1 Call for proposals

As explained above, we interleave resource scheduling and task execution. Thus, when a task needs to be allocated or a resource must be purchased, the agent in charge of its workflow is able to define the requirements for the next task. For example, the earliest starting time, defining the maximum task delivery time, the kind of resource that can develop the task, minimum skills (or licenses) required for performing the task, etc. Thus a task T^j is defined by a set of parameters $\langle pa_1^j, \dots, pa_m^j \rangle$ that defines the type of task which is going to be deployed and its characteristics (e.g. delivering a package from a depot to a workshop not earlier than 9:00 but not later than 17:00).

$$T^j = \langle pa_1^j, \dots, pa_m^j \rangle \quad (3.1)$$

Once the task to be allocated is defined, the agent a_0 in charge of the workflow starts an auction to allocate T_0^j . In a uni-attribute auction it would be assumed that the accomplishment of

T_0^j with any parameter configuration would produce the same utility for the auctioneer. However, in many domains, accomplishing the task with a different configuration of parameters will produce a different utility to the auctioneer. For example, finishing a task earlier may increase the auctioneer's utility, as it could start the next task earlier if it so wished. Thus, the auctioneer needs to specify not only the task T_0^j it wants to auction but also the set of attributes $AR_0 = ar_0^1, \dots, ar_0^n$ that influence its utility and which will be taken into account during the process of determining the auction winner for the bidders.

Therefore, when the auctioneer a_0 needs to allocate a task, it sends a call for proposals (CFP) to all the available bidders. This will include the definition of the task T_0^j which is going to be auctioned and the set of attributes AR_0 which will be taken into account during the allocation process:

$$\begin{aligned} CFP &= (T_0^j, AR_0) \\ CFP &= (T_0^j, \langle ar_0^1, \dots, ar_0^n \rangle) \end{aligned} \quad (3.2)$$

Additionally, the auctioneer might decide to make public the criteria that will determine the winner of the auction in order to help auction participants to bid optimally [15]. However, in our approach, if the auctioneer considers that making the WDP criteria public can compromise its privacy it can decide to keep this information secret.

After publishing the CFP, the auctioneer will leave a certain time for bidders to make their offers. After this period, the auctioneer will not accept more bids and will decide who the winner of the auction is.

Example 3.1. *The computer hardware manufacturing company CHM is producing a set of laptops. At a given point, it runs out of fans and needs to buy a bundle of 1,000 fans of a given size and material. CHM needs the fans to be delivered within 72 hours; however, they would appreciate delivery of the fans earlier as this would give them a higher flexibility when building the laptops. Moreover, CHM has an strict environmental policy and cannot exceed a carbon footprint of more than 190kg of CO₂ per laptop. Thus, it requires a certification proving that the carbon footprint of each fan is not higher than 3 CO₂ kg, moreover, CHM would give a higher value to fans with a lower footprint as this would allow the company to invest this CO₂ quantity into more critical components.*

Using PUMAA the agent in charge of manufacturing will summon an auction defining a task T_{chm}^{fan} which will be defined by the type of task ('manufacturing and delivering 1,000 fans');

the characteristics of the fans (size, material, working voltages, etc.); the maximum delivery time (72h) and the maximum acceptable carbon footprint (3kg):

$$T_{chm}^{fan} = \langle \text{'build 1,000 fans', 'fan characteristics', 'max delivery time = 72h', 'max CO}_2=3\text{kg'} \rangle \quad (3.3)$$

Moreover, the auctioneer has to specify that besides the cost, the delivery time offered (dt) and carbon footprint (cf) will be taken into account during the winner determination:

$$AR_0 = \langle dt, cf \rangle \quad (3.4)$$

$$CFP = \left(T_{chm}^{fan}, \langle dt, cf \rangle \right) \quad (3.5)$$

3.2.2 Bidding

The main goal of a bidder is to maximize its utility, to do so, it needs to win auctions to perform tasks and try to maximize its profit (the difference between the payment it receives for a task and the economic cost it involves).

Agents which are present in the market will receive the call for proposals from the auctioneers. Every time an agent a_i receives a CFP for performing a task T_0^j with a set of attributes AR_0 it evaluates if it can perform the desired task (if it meets the described requirements, if it has enough capacity for performing that task or if it prefers to wait for another auction).

If a_i decides to participate in the auction, it has to provide a bid B_i containing the set of attributes $AT_i = at_i^0, \dots, at_i^n$ which define the attributes required in AR_0 (note that $|AT_i| = |AR_0|$) and the economic cost b_i for performing the task with the given attributes (Equation 3.6). In the case that a bidder does not know which of its possible bids has higher chances of winning the auction it may decide to send more than one bid with different attribute configurations (e.g. an normal price and quality in one bid and a better quality at a higher price in another bid) so the auctioneer decides which is the best option.

$$B_i = (b_i, AT_i) \quad (3.6)$$

The attributes AT_i which the bidder offers in the bid may correspond or not to the true value of attributes that it pretends to deliver if it wins the task (AT_i^t) (note that AT_i^t does not necessary have to correspond to the final attributes delivered AT_i' , as the bidder may experience

unexpected incidents during the task performance or the bidder may intentionally lie). Given the task T_0^j defined in the CFP and the attributes AT_i^t , a_i has an economic cost b^t for performing the task. This economic cost can be defined as $b_i^t = v_i(T_0^j, AT_i^t)$. As happens with the values of the attributes, the bidder can offer its real economic cost b_i^t (which corresponds to the value it gives to performing the task with a given set of attributes) or a different value depending on what the bidder believes will maximize its utility:

$$\begin{aligned} u_i(p_i, b_i^t) &= p_i - b_i^t \\ u_i(p_i, T_0^j, AT_i^t) &= p_i - v_i(T_0^j, AT_i^t) \end{aligned} \quad (3.7)$$

In an incentive compatible mechanism, agents bid truthfully. Meaning that agents provide their true valuation in bids: $B_i = (b_i^t, AT_i^t)$. In a non-incentive compatible mechanism agents may obtain higher utilities by providing false values (either respect to the price or to the attributes):

$$B_i = (b_i, AT_i) | b_i \neq b_i^t \vee AT_i \neq AT_i^t \quad (3.8)$$

For example, in a false bid the bidder could offer a delivery time lower than the one it intends to provide in order to beat a bidder with a better offer.

Example 3.2. *Three companies are interested in the CFP sent by CHM: the FastCompany (FC), the SlowCompany (SC) and the UntrustableCompany (UC). The first one is a reliable but expensive company which can perform the required task in 50 hours and producing 3CO₂ kg per fan with a cost of 10,000€. SC is a cheaper company which has slower production methods, thus it can perform the required task in 70 hours with a carbon footprint of 2.9 CO₂ kg with a cost of 8,000€. Finally, UC is a company which announces a misleading offers at inflated prices in order to obtain new customers; its true values are 72 hours, 3CO₂ kg and a cost of 8,950€.*

Given that FC and SC are honest companies their bids correspond to their true values. For FC its economic cost is $b_{FC} = v_{FC}(T_{chm}^{fan}, 50, 3) = 10,000$ and, consequently, the bid it will provide is $B_{FC} = (10,000, (50, 3))$. The cost for SC is $b_{SC} = v_{SC}(T_{chm}^{fan}, 70, 2.9) = 8,000$, thus, its bid will be $B_{SC} = (8,000, (70, 2.9))$.

On the other hand, UC is a greedy company which desires to increase its benefits with dubious strategies. Despite its real cost being $b_{UC} = v_{UC}(T_{chm}^{fan}, 72, 3) = 8950$ it will ask for a higher price: 9,000€. Equally, it announces that it will finish sooner than its real delivery date, in order to try to win the auction. Thus, its bid could be $B_{UC} = (9,000, (65, 3))$.

3.2.3 Winner determination

Once the period for submitting bids expires, the auctioneer filters all the feasible bids (bids which fulfill the requirements presented in the CFP) and, in case of receiving more than one bid from any participant, the auctioneer selects only the best offer submitted by that participant. Then, it executes the auction upon deciding who the winner of the auction is.

In a forward single-attribute auction the bids evaluation is implicit (the higher the bid, the better): $V_0(b_i) = b_i$. Thus, the auctioneer selects the allocation which maximizes the value of V_0 .

$$winner = argmax_i(V_0(b_i)) \quad (3.9)$$

In the same way, in reverse auctions, the winner of the auction is the one with the lowest $V_0(b_i)$:

$$winner = argmin_i(V_0(b_i)) \quad (3.10)$$

Given that in multi-attribute auctions bids are composed of more than one attribute, we propose that bids are evaluated using a multi-criteria aggregation function [32] to combine the bid price b_i with the bundle of attributes AT_i :

$$V_0(b_i, AT_i) \quad (3.11)$$

This aggregation function acts as evaluation function $V_0(b_i, AT_i)$ (similarly to Che's Score function [15]) and must be defined according to the auctioneer's expected utility function. In a multi-attribute domain, the utility function (Equation 3.12) of an auctioneer agent can be defined as the valuation v_0 the auctioneer gives to the auctioned item (in the present case, a task) minus a function f_0 . f_0 defines the auctioneer's internal valuation of the payment it makes for the task and how it evaluates the delivered attributes of the task AT'_i .

$$u_0(T_0^j, p_i, AT'_i) = v_0(T_0^j) - f_0(p_i, AT'_i) \quad (3.12)$$

where $v_0(T_0^j)$ is a function which describes the value which a_0 gives to T_0^j and $f_0(p_i, AT'_i)$ the value which gives to the payment and the bundle of delivered attributes.

Consequently, the expected utility function (Equation 3.13) will be the same function but using the economic value of the bid b_i instead of the payment p_i and the bided attributes instead of the ones delivered.

$$\bar{u}_0(T_0^j, b_i, AT_i) = v_0(T_0^j) - f_0(b_i, AT_i) \quad (3.13)$$

Assuming that all the bids the auctioneer evaluates are feasible (meet the requirements of T_0^j), the best way to choose the winning offer is to minimize f_0 . Therefore, the evaluation function V_0 that will determine the winner of the auction should be equal to f_0 . However, for a function to be used as evaluation function it must satisfy certain requirements: it must be monotonic, continuous, real-valued and bijective (these three aspects are discussed in more detail in Section 3.3). Therefore, under certain circumstances, f_0 might not be valid to be used as evaluation function. In such scenario, the mechanism designer should pick an aggregation function V_0 as similar to f_0 as possible in order to determine the winner of the auction which will maximize the auctioneer's expected utility:

$$V_0(b_i, AT_i) \approx f_0(b_i, AT_i) \quad (3.14)$$

Thus, as in single-attribute reverse auctions, the winning bid will be the one minimizing the output of V_0 .

$$\operatorname{argmin}_i (V_0(b_i, AT_i)) \quad (3.15)$$

with $V_0(AT_i) \in \mathfrak{R}$.

It is important to note that the number of input variables of V_0 will correspond to the number of attributes to be considered plus the economic cost ($|AR_0| + 1$). Therefore, an auction where the cost, the delivery time and a quality parameter are used has an evaluation function with three input variables: $V_0(a, b, c)$.

In case of a draw where two bids share the best evaluation, the tie must be broken using an arbitrary tie-break rule defined by the auctioneer [52]. In this situation, given the Vickrey nature of the mechanism, both the winner of the auction and the non-winning bidder will obtain 0 payoff as they will receive the exact amount they asked for.

Despite this mechanism has been designed for considering numerical attributes, if the preferences of the auctioneer are ordered, PUMAA can also admit qualitative attributes. For that purpose the auctioneer needs to map the possible values of the attribute according to its preferences [76] (for example, using a function which assigns a numerical value to each possible domain item). In this way, the qualitative attributes can be treated as quantitative attributes. However, if the auctioneer's preferences cannot be linearly expressed (partially-ordered preferences), qualitative attributes cannot be handled. In such case, we recommend to use preference-based auctions such as Bellosta et al.'s approach [7].

Example 3.3. CHM values receiving all the fans on time for 10,000€ ($v_0(T_{chm}^{fan}) = 10,000$). Moreover, it values each extra hour it can obtain due to an earlier delivery at 50€ and each CO₂

	b	dt	cf	$V_{CHM}(b_i, dt_i, cf_i)$	$\bar{u}_{CHM}(T_{chm}^{fan}, b_i, dt_i, cf_i)$
HC	10,000	50	3.0	8,900	1,100
SC	8,000	70	2.9	7,800	2,200
UC	9,000	65	3.0	8,650	1,350

Table 3.1: Valuations of the bids received by CHM and the expected utility \bar{u} (the utility they would produce if the task is delivered as agreed, and the payment is the one bid)

kilogram saved per fan as 1€ . Thus, the utility of CHM is the one defined in as follows:

$$u_{CHM}(T_{chm}^{fan}, p, dt'_i, cf'_i) = 10,000 - p + (72 - dt'_i) * 50 + (3 - cf'_i) * 1,000. \quad (3.16)$$

In consequence, the evaluation $V_{chm}(b_i, dt_i, cf_i)$ function it uses to determine the winner of the auction is the following:

$$V_{CHM}(b_i, dt_i, cf_i) = b_i - (72 - dt) * 50 - (3 - cf) * 1,000. \quad (3.17)$$

Given the bids CHM receives (Table 3.1) the winner of the auction is the Slow Company with a bid evaluation of 7,800. It is important to note that if the SC had not participated in the auction, the winner would have been UC due to its false bid (if it had bid truthfully it would have obtained a valuation of 8,950, being the worst bid).

3.2.4 Payment

In Vickrey uni-attribute auctions, the payment of the winner corresponds to the economic amount offered by the second best bid. Regarding multi-attribute auctions, Che [15] states that second price or Vickrey payment must consider all the attributes involved in the decision process. Particularly, he states that the auction winner can provide any set of attributes that equals the evaluation obtained by the second best bid; however, in the manufacturing and the supply chain domains, the auctioneers need bidders to stick to the offers they submitted in order to plan future tasks. This poses the need of adapting the standard second price payment to the particularities of the domain. Taking this into account, we develop our payment mechanism which is inspired in Google Position auctions [84]. Those auctions deal with fixed attribute values (which cannot vary between the bidding process and the payment) and variable payments. However, conversely to Google's approach, we have to concern about bidders respecting the attributes they offered.

In our proposal, the auctioneer fixes the bid attributes provided by the winner (the auctioneer expects that the delivered task attributes are, at least, as good as in the ones in the

winner bid) and pays the winner the amount just necessary to beat the second highest bid (Equation 3.18). In other words, the payment p the winner receives is the price it should have bid to obtain the same evaluation as the second highest bid¹.

$$V_0(p_1, AT_1) = V_0(b_2, AT_2) \quad (3.18)$$

Where p_1 is the payment of the single winner in PUMAA, AT_1 the attributes of the winner bid, and b_2, AT_2 the components of the second best bid.

This strategy does not prevent the bidders from lying regarding their attributes, as including a false attribute could increase the chances of winning the auction whilst not being penalized in the payment. For example, a bidder could submit a bid saying that it will finish its task in 10 minutes when it would actually finish the task in 15 minutes. This lie would have increased the chances of the bidder winning the auction.

Thus we adapt the payment mechanism in order to minimize the impact of poorly performed tasks and to penalize dishonest bidders²: e.g. when an agent delivers a set of attributes AT'_i worse than the attributes AT_i offered in the bid. When a bidder lies to win the auction, given a set of delivered attributes AT'_i , payment is obtained by computing how much the winner should have bid to obtain the winning evaluation but with the delivered set of attributes. As the payment is performed after the execution of a task, the auctioneer can measure the quality of the results and observe that instead of obtaining AT_i , it got AT'_i . When this situation arises, the payment mechanism is based on preserving the valuation of the allocation (trying to obtain the same valuation with the payment and the delivered attributes as the original winning bid).

$$V_0(p_1, AT'_1) = V_0(b_1, AT_1) \quad (3.19)$$

Where b_1 is the price offered by the winner bid and AT'_1 the attributes delivered by the auction winner.

Therefore, if we define $V_0^{-1}(x_i, AT_i) = b_i$ as the partial inverse function of $V_0(b_i, AT_i) = x_i$ which, given a set of attributes AT_i and the result x_i of an evaluation, returns the economic amount of the bid b_i , we can define the payment function in \mathfrak{R} according to Equation 3.20.

$$p_1 = \begin{cases} V_0^{-1}(V_0(b_2, AT_2), AT_1) & \text{if } AT'_1 \succeq AT_1 \\ V_0^{-1}(V_0(b_1, AT_1), AT') & \text{if } AT'_1 \prec AT_1 \end{cases} \quad (3.20)$$

where operator \succeq means *the same or better than* and \prec *worse than*.

¹For the sake of simplicity, we consider that bids are ranked and b_1 corresponds to the best bid and b_2 to the second best bid

²We assume that a bidder is capable of accurately estimate its attributes

Therefore, the utility derived for the auctioneer once the auctioned task is finished is as follows:

$$u_0 = \begin{cases} v_0(T_0^j) - f_0(V_0^{-1}(V_0(b_2, AT_2), AT_1), AT_1') & \text{if } AT_1' \geq AT_1 \\ v_0(T_0^j) - f_0(V_0^{-1}(V_0(b_1, AT_1), AT_1'), AT_1') & \text{if } AT_1' < AT_1 \end{cases} \quad (3.21)$$

Obviously, the definition of $V_0(b_i, AT_i)$ also conditions the payment since an inappropriate evaluation function may preclude the payment calculation. In addition, V_0 also affects the economic amount that bidders who fail to deliver their tasks receive: depending on V_0 , a bidder who largely fails to deliver the agreed attributes can see its payment reduced, it may receive no payment (a payment value of 0) or it may have to pay a fee (receive a negative payment). For instance, when using multiplicative functions like the product, a bidder can only receive a payment reduction (as the worse is the attribute, the smaller the payment but being always higher than 0); on the other hand, when using additive functions like the sum or the weighted sum, if the bidder provides a very bad set of attributes it might receive a zero or negative payment.

With this payment rule bidders are encouraged to bid truthfully. When the winning bidder delivers its product or service as agreed, the payment it receives is based upon the second best bid (preventing mechanism manipulation from the winner side): offering a lower amount would result in the same payment but increases the chances of working under the production cost; offering a higher amount would result in the same payment and payoff but risks winning the auction. Bidding attributes worse than what can actually be delivered also increases the risk of not winning the auction without increasing the winner's payoff. Finally, delivering the product under conditions poorer than those agreed produces a payment and payoff reduction. The incentive compatibility discussion regarding the mechanism is widely extended in Section 3.4.

This two case payment mechanism also reduces the utility loss an auctioneer can suffer when a bidder delivers an item under worse conditions than those agreed, whether this situation is caused intentionally or accidentally. Particularly, the amount the auctioneer pays in this case is the amount needed to equal the real utility with the expected utility maximized during the WDP. Consequently, the output of f_0 when the bidder delivers a bad set of attributes is equal or lower than the output of f_0 when the bidder delivers the agreed attributes. In this way, the utility of the auctioneer is preserved even if the bidder incorrectly reported the values of their attributes AT_i .

Example 3.4. *Continuing with the previous example, consider that the winner of the auction, SC, delivers the fans as agreed (in 70 hours and with a carbon footprint of 2.9CO₂ kg per fan.)*

Given that SC fulfilled the bid agreement, its payment corresponds to the economic amount it should have offered to equal the second best bid provided by UC (with a valuation of 8,650):

$$\begin{aligned} V_0(p_{sc}, AT_{sc}) &= V_0(b_{uc}, AT_{uc}) \\ p_{sc} - (72 - 70) * 50 - (3 - 2.9) * 1,000 &= 9,000 - (72 - 65) * 50 - (3 - 3) * 1,000 \quad (3.22) \\ p_{sc} &= 8,650 + 200 = 8,850 \end{aligned}$$

Thus, the payment SC will receive is 8,850€, the payoff or utility which SC derives from the work is 850 (8,850-8,000) and the utility of the auctioneer CHM :

$$u_{chm}(T_{chm}^{fan}, p_{sc}, dt'_{SC}, cf'_{SC}) = 10,000 - 8,850 + (72 - 70) * 50 + (3 - 2.9) * 100 = 1,260. \quad (3.23)$$

Example 3.5. *The winner of the auction, SC fails to deliver fans as agreed: despite delivering the fans in 70 hours, instead of providing a carbon footprint of 2.9CO₂ kg per fan it provides a carbon footprint of 3CO₂ kg per fan.*

Given that SC had not been able to perform the task, its payment corresponds to the amount it should have bid to obtain the same evaluation with $AT'_{SC} = (8850, 70, 3)$ than with the original bid ($AT_{SC} = (8850, 70, 2.9)$):

$$\begin{aligned} V_0(p_{sc}, AT'_{sc}) &= V_0(b_{sc}, AT_{sc}) \\ p_{sc} - (72 - 70) * 50 - (3 - 3) * 1,000 &= 8,000 - (72 - 70) * 50 - (3 - 2.9) * 1,000 \quad (3.24) \\ p_{sc} &= 7,800 + 100 = 7,900 \end{aligned}$$

Thus, the payment SC will receive is 7,900€, the payoff which SC derives from the work is negative (7,900-8,000=-100) and the utility of the auctioneer CHM is the following:

$$u_{chm}(T_{chm}^{fan}, p_{sc}, dt'_{SC}, cf'_{SC}) = 10,000 - 7,900 + (72 - 70) * 50 + (3 - 3) * 100 = 2,200. \quad (3.25)$$

In this case it can be seen that, with PUMAA, providing a false attribute can compromise the benefits which a bidder obtains from winning the auction whilst the auctioneer is not harmed (in this case it may even obtain a higher benefit due to the compensation for the agreement breach).

3.3 Evaluation Function Characteristics

In the previous section we stated that the definition of the multi-criteria function which acts as the evaluation functions is a key aspect for the mechanism development. The evaluation func-

tion is used to decide the auction winner but it also conditions the computation of payment. Therefore, a bad decision in the evaluation function may invalidate the mechanism because an inappropriate evaluation can preclude the computation of payment.

In this section we first define the requirements a multi-criteria function must fulfill in order to be used as evaluation function. Then, we present a few examples of functions which can be used within PUMAA. The study made in this section is intended specifically for PUMAA, however, the results may be extended to other second-price, score-based mechanisms such as Che's second-offer auctions [15] or Porter's fault-tolerant auction [65].

3.3.1 Multicriteria function as evaluation function: requirements

In order to use a multicriteria function as an evaluation function, it must fulfill a set of conditions within the range of the attributes. First, the functions used for the evaluation must return a real number so that the different bids can be analytically compared and ranked from the best to the worst; then, the functions must be monotonic, giving a better valuation for a better bid. So the payment can be calculated, the evaluation function must be bijective for the price attribute.

- **Real-valued Function:** Given a set of bids, the evaluation function must return a real number evaluation for each bid so that the bids can be ranked and compared. As the auction payment involves the valuation obtained by the second best bidder and does not directly correspond to the price bid by the winner, multi-criteria methods which result in ranked lists without a numeric rating for each item cannot be used, as the payment cannot be calculated.
- **Monotonicity:** The evaluation function must be monotonic. If one of the attributes of a bid is improved the result of the evaluation function will change, consequently, granting that a better bid will not obtain a worse evaluation. This property also implies that, for every possible value inside the domain attribute, the evaluation function will return a value. It is important to observe that the monotonicity requirement is applied only to the range of values that an attribute can take, allowing functions which are only monotonic in the attribute range to be used as evaluation functions. For example in a situation where all the attributes take values which are in the positive number domain, the Euclidean norm could be used as evaluation function.
- **Bijection:** In order to allow the mechanism to calculate the payment, the evaluation function must have a bijective behavior regarding the price attribute. This means that,




Figure 3.3: Examples of functions that cannot be used as evaluation functions and their possible alternatives.

given the bid attributes values and the result of the evaluation function, the cost attribute of the bid can take only one possible value. If this condition is not fulfilled, then, the auctioneer would be unable to calculate the payment that the winner should receive as Equation 3.20 would have more than one solution. In other words, given the function $V_0(b_i, AT_i) = x_i$, the inverse function will be $V_0^{-1}(x_i, AT_i) = b_i$ where b_i can take just one value. In addition, to avoid the payment having an infinite value, the inverse function V_0^{-1} can not have any vertical asymptote within the non-economic attribute domain.

Figure 3.3 illustrates examples of functions which do not commit the monotonicity and the bijection requirements: the left chart shows a uni-dimensional example where only the economic cost is considered ($f_0(b_i) = b_i - 2\sin(b_i)$) and the right one presents a two-dimensional case where the economic price and another attribute are considered ($f_0(b_i, at_i) = at_i + b_i - 5\sin(at_i)$). In neither of the two cases f_0 can be used as evaluation function as it is not monotonic nor bijective. Therefore, the mechanism designer needs to choose an evaluation function V_0 with similar outputs to f_0 whilst assuring that it commits the three above stated requirements. For instance, for the first case it could use $V_0(b_i) = b_i - 0.5$ whilst in the second one a suitable solution could be $V_0(b_i, at_i) = at_i + b_i - 1$.

These conditions may seem very restrictive, however, it must be stated that these requirements are limited to the values that attributes can take. For example imagine an auction where two attributes are involved: economic cost b and delivery time dt . In this scenario it is reasonable to assume that both the delivery times and the economic cost which bidders offer belong

to the domain of the positive real numbers (\mathfrak{R}^+): To the best of our knowledge science still does not allow time travel so every task will be delivered at a future timepoint ($dt > 0$), moreover, an agent will not pay to work therefore, its minimum bid will be 0 (to work for free). With this information, the auctioneer could use an aggregation function such as this:

$$V_0(b_i, dt_i) = (b_i)^2 + (dt_i)^2 \quad (3.26)$$

Despite the fact that $V_0(b_i, dt_i)$ is not monotonic or bijective for all the \mathfrak{R} , it commits to the three requirements when only the (\mathfrak{R}^+) are taken into account. Thus, it is a suitable evaluation function for the hypothetical case we have just described.

3.3.2 Multicriteria function as evaluation function: examples

There are a wide range of functions which, given a well defined domain, can act as evaluation functions. In this section we present some examples assuming that all the attribute values belong to the domain of positive real numbers and are normalized.

Product and weighed sum

The product³ (Equation 3.27) and the weighed sum (Equation 3.29) can be used as evaluation functions. Due to the simplicity of these functions, the payment functions which are derived from them are also simple and practical to use (Equations 3.28 and 3.30).

Product:

$$V_0(b_i, AT_i) = b_i * \prod_{j=1}^n at_i^j \quad (3.27)$$

$$V_0(p_1, AT'_1) = \begin{cases} V_0(b_2, AT_2) & \text{if } AT'_1 \geq AT_1 \\ V_0(b_1, AT_1) & \text{if } AT'_1 < AT_1 \end{cases}$$

$$p_1 * \prod_{j=1}^n at_i^{j'} = \begin{cases} b_2 * \prod_{j=1}^n at_2^j & \text{if } AT'_1 \geq AT_1 \\ b_1 * \prod_{j=1}^n at_1^j & \text{if } AT'_1 < AT_1 \end{cases}$$

$$p_1 = \begin{cases} \frac{b_2 * \prod_{j=1}^{j=n} at_2^j}{\prod_{j=1}^n at_1^j} & \text{if } AT'_1 \geq AT_1 \\ \frac{b_1 * \prod_{j=1}^{j=n} at_1^{j'}}{\prod_{j=1}^n at_1^j} & \text{if } AT'_1 < AT_1 \end{cases} \quad (3.28)$$

³The product can only be used as evaluation function if the non-economic attribute domains do not contain number 0. Otherwise the payment could be infinite.

Weighted sum:

$$V_0(b_i, AT_i) = \mu_0 b_i + \sum_{j=1}^n \mu_j a t_i^j \quad (3.29)$$

$$V_0(p_1, AT'_1) = \begin{cases} V_0(b_2, AT_2) & \text{if } AT'_1 \geq AT_1 \\ V_0(b_1, AT_1) & \text{if } AT'_1 < AT_1 \end{cases}$$

$$\mu_0 p_1 + \sum_{j=1}^n \mu_j a t_1^{j'} = \begin{cases} \mu_0 b_2 + \sum_{j=1}^n \mu_j a t_2^j & \text{if } AT'_1 \geq AT_1 \\ \mu_0 b_1 + \sum_{j=1}^n \mu_j a t_1^j & \text{if } AT'_1 < AT_1 \end{cases}$$

$$p_1 = \begin{cases} \frac{\mu_0 b_2 + \sum_{j=1}^n \mu_j (a t_2^j - a t_1^j)}{\mu_0} & \text{if } AT_1 \geq AT_1^v \\ \frac{\mu_0 b_1 + \sum_{j=1}^n \mu_j (a t_1^{j'} - a t_1^j)}{\mu_0} & \text{if } AT_1 < AT_1^v \end{cases} \quad (3.30)$$

where $\mu^j \in [0, 1]$ is the weight of each summation term and $\mu_0 + \sum_{j=1}^n \mu_j = 1$.

Mathematical norms

Another example of possible evaluation functions are certain mathematical norms. For example assuming that the attribute domain belongs to the set of positive numbers plus 0, the Euclidean norm (Equation 3.31) can be used.

In contrast to the product, this evaluation function would favor bids with more polarized attributes (see Section 6.3.3). The corresponding payment function is given by Equation 3.32.

$$V_0(b_i, AT_i) = \sqrt{b_i^2 + \sum_{j=1}^n (a t_i^j)^2} \quad (3.31)$$

$$V_0(p_1, AT'_1) = \begin{cases} V_0(b_2, AT_2) & \text{if } AT'_1 \geq AT_1 \\ V_0(b_1, AT_1) & \text{if } AT'_1 < AT_1 \end{cases}$$

$$\sqrt{p_1^2 + \sum_{j=1}^n (a t_1^{j'})^2} = \begin{cases} \sqrt{b_2^2 + \sum_{j=1}^n (a t_2^j)^2} & \text{if } AT'_1 \geq AT_1 \\ \sqrt{b_1^2 + \sum_{j=1}^n (a t_1^j)^2} & \text{if } AT'_1 < AT_1 \end{cases}$$

$$p_1 = \begin{cases} \frac{\sqrt{b_2^2 + \sum_{j=1}^n ((a t_2^j)^2 - (a t_1^j)^2)}}{\sqrt{b_1^2 + \sum_{j=1}^n ((a t_1^{j'})^2 - (a t_1^j)^2)}} & \text{if } AT'_1 \geq AT_1 \\ \frac{\sqrt{b_1^2 + \sum_{j=1}^n ((a t_1^{j'})^2 - (a t_1^j)^2)}}{\sqrt{b_1^2 + \sum_{j=1}^n ((a t_1^{j'})^2 - (a t_1^j)^2)}} & \text{if } AT'_1 < AT_1 \end{cases} \quad (3.32)$$

It is important to remark that not all the norms can be used as evaluation function. For example the Chebyshev norm [13] cannot be used as the evaluation function as it is not a bijective function and the payment could not be calculated.

Weighted sum of functions

All the functions commented upon above treat the different attributes in the same way, however, in some domains, the attributes require individual treatment and modeling. For these cases, a multicriteria function which fits the evaluation function requirements is the weighted sum of functions (WSF). Equation 3.33 shows that each attribute is individually evaluated using a function $g_j(at_i^j)$ and the different results are then aggregated using a weighted function. This multicriteria method gives the advantage of being highly adaptable to the domain, however, in order to be used as evaluation function all the functions $g_j(at_i^j)$ must fulfill the requirements presented in Section 3.3.1. Moreover, the payment function would depend upon the attribute functions $g_j(at_i^j)$ as well as their inverse functions.

$$V_0(b_i, AT_i) = \mu_0 g_0(b_i) + \sum_{j=1}^n \mu_j g_j(at_i^j) \quad (3.33)$$

$$V_0(p_1, AT_1') = \begin{cases} V_0(b_2, AT_2) & \text{if } AT_1' \succeq AT_1 \\ V_0(b_1, AT_1) & \text{if } AT_1' \prec AT_1 \end{cases}$$

$$\begin{aligned} \mu_0 g_0(p_1) + \sum_{j=1}^n \mu_j g_j(at_1^{j'}) &= \begin{cases} \mu_0 g_0(b_1) + \sum_{j=1}^n \mu_j g_j(at_1^j) & \text{if } AT_1' \succeq AT_1 \\ \mu_0 g_0(b_2) + \sum_{j=1}^n \mu_j g_j(at_2^j) & \text{if } AT_1' \prec AT_1 \end{cases} \\ p_1 &= \begin{cases} \frac{g'_0(\mu_0 g_0(b_2) + \sum_{j=1}^n (\mu_j g_j(at_2^j) - \mu_j g_j(at_1^j)))}{\mu_0} & \text{if } AT_1' \succeq AT_1 \\ \frac{g'_0(\mu_0 g_0(b_1) + \sum_{j=1}^n (\mu_j g_j(at_1^{j'}) - \mu_j g_j(at_1^j)))}{\mu_0} & \text{if } AT_1' \prec AT_1 \end{cases} \end{aligned} \quad (3.34)$$

where $g_0(b_i)$ and $g_j(at_i^j)$ are the functions which evaluate the attributes, $g'_0(x)$ and $g'_j(x)$ their inverse functions, and μ_j are the weights of each attribute function.

3.4 Incentive Compatibility

One of the desired properties of any new mechanism, such as PUMAA, is to be incentive compatible. In this section we discuss and prove the incentive compatibility of PUMAA when

assuming that there are no externalities (bidders are just focused on the present auction without taking into account if there will be more auctions and only use information which can be extracted from the auction itself).

Incentive compatibility can be analyzed with a case study of the utility which bidders can derive from their strategies, as follows:

- **Strategy 1:** Offering bids with better attributes than the ones which a bidder can afford would provoke a payment reduction (as shown in Equation 3.20). In particular the bidder would receive the amount it would have had to offer to win the auction with the delivered attributes. Thus, providing better attributes than the real ones can at most provide the same payoff as truthful bidding (TB). Moreover this strategy may derive negative payoff utilities whilst TB cannot.
- **Strategy 2:** Offering worse attributes than the ones which the bidder intends to deliver reduces the chances of winning the auction, increasing the chances of obtaining a utility of 0 (the worst utility than can be obtained with TB).
- **Strategy 3:** Offering an economic cost lower than the real value (underbidding) increases the chances of winning the auction, however it also increases the chances of working under the bidder's cost (providing a negative payoff). If a bidder obtains a positive payoff when underbidding it is because the second best bid has a higher valuation than the winning bidder's true value; in this case, the bidder would have obtained the same payoff using TB.
- **Strategy 4:** Offering a higher economic cost than the real one (overbidding) reduces the chances of winning the auction (obtaining 0 utility). In the case that a bidder obtains a payoff by overbidding, the payoff will be the same it would have obtained by TB as its payment is conditioned by the second best bid valuation, not by the winning one (its bid).
- **Strategy 5:** In the case of a draw, an arbitrary tie-break rule decides the winner. In this situation, given the Vickrey-based payment mechanism, if both bidders bid truthfully they will obtain the same payoff (0) [52]. If one of them lies in order to obtain a lower valuation (and win the auction) it will also obtain a payoff of 0 (as its payment will correspond to the valuation of the second best bid which it is actually the same true valuation of the winner).

PUMAA appears to be strategy proof, however this informal analysis is not enough to demonstrate its incentive compatibility.

There are two main ways of demonstrating the incentive compatibility of a mechanism, the first one is to demonstrate that the mechanism commits the sufficient conditions of exactness, participation, monotonicity and criticality provided by Lehman et al. [46]; the second way is to demonstrate that there is no situation where a bidder obtains a higher utility by lying in its bids than by providing truthful bids.

In this section we first describe how PUMAA fulfills the properties described by Lehman et al. Then, we demonstrate that there is no case where a bidder obtains higher utility by lying than by bidding truthfully. In order to do so, using a constraint solver, we model an inequation system defining the auction mechanism and the case where a bidder can obtain higher utilities by lying. If this system is consistent, this will mean that the mechanism is not incentive compatible.

3.4.1 Condition satisfaction

According to Lehman et al. [46] a mechanism is incentive compatible if all these sufficient conditions are satisfied:

- **Exactness** postulates that a single minded bidder receives exactly the set of goods it desires or nothing. In single-item auctions this means that the bidder gets the item it bid for or nothing.
- **Participation** deems that unsatisfied bidders pay zero and their utility is zero. In other words, a non-participating or non-winning bidder neither pays nor receives any money.
- **Monotonicity** requires that if a bidder increases its bid (decreases in reverse auctions) the bidder still wins the auction.
- **Criticality** claims that each winning bidder pays (receives in reverse auctions) the lowest value it could have declared and still be allocated the goods it requested.

PUMAA meets all those requirements:

- **Exactness:** In PUMAA the auctioned item is assigned to only one winner, moreover, the only way of winning the auctioned item is to participate in the auction offering the best bid. In consequence, bidders cannot obtain items for which they have not bid: they get the item for which they have bid or nothing.

- **Participation:** In our proposal there is no fee to access the auction. Thus, if an agent does not win, it does not pay anything obtaining 0 utility.
- **Monotonicity:** this property is strictly related to the evaluation function. When using a monotonic function as an evaluation function (e.g. the sum, the product or the weighted sum) the act of improving any of the attributes causes the bidder to gain a higher evaluation (see equations below). In consequence, improving a winning bid cannot cause a bidder to lose the auction.

Below (Equations 3.35 to 3.37) this situation is illustrated when using the sum, the product and the weighted sum as evaluation functions. Note that the lowest value of V_0 is the one which wins the auction and that when we modify one of the attributes (adding a -1) we are improving it.

Sum:

$$\begin{aligned} V_0(b_i, AT_i) &= b_i + \sum_{j=1}^n at_i^j \\ (b_i - 1) + \sum_{j=1}^n at_i^j &< b_i + \sum_{j=1}^n at_i^j \\ b_i + \sum_{j=1}^n (at_i^j - 1) &< b_i + \sum_{j=1}^n at_i^j \end{aligned} \quad (3.35)$$

Product:

$$\begin{aligned} V_0(b_i, AT_i) &= b_i + \prod_{j=1}^n at_i^j \\ (b_i - 1) + \prod_{j=1}^n at_i^j &< b_i + \prod_{j=1}^n at_i^j \\ b_i + \prod_{j=1}^n (at_i^j - 1) &< b_i + \prod_{j=1}^n at_i^j \end{aligned} \quad (3.36)$$

Weighted sum:

$$\begin{aligned} V_0(b_i, AT_i) &= \mu_0 b_i + \sum_{j=1}^n \mu_j at_i^j \\ \mu_0 (b_i - 1) + \sum_{j=1}^n \mu_j at_i^j &< \mu_0 b_i + \sum_{j=1}^n \mu_j at_i^j \\ \mu_0 b_i + \sum_{j=1}^n \mu_j (at_i^j - 1) &< \mu_0 b_i + \sum_{j=1}^n \mu_j at_i^j \end{aligned} \quad (3.37)$$

- **Criticality:** The payment mechanism computes the payment by matching the evaluation of the second best bid with the evaluation of the payment and the attributes of the winning bid ($V_0(b_2, AT_2) = V_0(p, AT_1)$). This ensures that the winning bidder will receive the amount it should have tendered to obtain the same evaluation as the second best bid, ensuring that the payment is the minimum amount that was needed to win the auction. This is given that the with the bidded attributes are equal, thereby respecting criticality.

Since PUMAA has the four properties mentioned above, we can deduce that the mechanism is incentive compatible.

3.4.2 Proof by using a solver

An alternative method for testing incentive compatibility is to analyze the utility which a bidder derives from a given strategy using a constraint solver [64]. To demonstrate that truthful bidding is the dominant strategy it has to be proven that, for any feasible bid, the utility of a bidder is higher or equal when bidding truthfully than when providing false attributes ($AT_i \neq AT_i^t$) or an economic bid different from the true value ($b_i \neq b_i^t$):

$$\forall i \in \mathbb{N}, \forall (b_i, b_i^t, AT_i, AT_i^t) \in \mathbb{R} > 0 : \{(u_i(b_i, AT_i, p) \leq u_i(b_i^t, AT_i^t, p')) | (b_i^t \neq b_i) \vee (AT_i \neq AT_i^t)\} \quad (3.38)$$

where $u_i(b, AT, p)$ corresponds to the utility function of the bidder i , b_i^t and AT_i^t are its true values, b_i and AT_i are the bid values, p the payment the bidder will receive and p' the payment it would have received bidding truthfully.

Therefore, if we model the auction mechanism as a constraint satisfaction problem we can try to find a counterexample which contradicts Equation 3.38, a case where the utility of a bidder would have been higher when lying than when telling the truth. If this case exists, then the mechanism is not incentive compatible. Thus, our goal is to find out whether these cases exist or not.

Problem definition

The auction mechanism and the counter-example can be modeled as an inequation system. If a constraint solver [4] is able to find a solution for the inequation system it will show that there exists, at least, one case in which the utility of the bidder is higher when lying than when bidding honestly. This would refute the hypothesis that the mechanism is incentive compatible. It is important to take into account the kind of functions which define the auctioneer utilities, the WDP and the payment rule as this if the solver must support linear or non-linear arithmetic.

The result of this experiment may vary depending of the evaluation function used for the auction mechanism. Therefore, the results of this test indicate if the mechanism is incentive compatible given a determinate evaluation function. Particularly, we test the incentive compatibility regarding the sum, the product and the weighted sum. Since these aggregation functions can accept from 1 to n attributes, we will simplify the notation using directly the result of a function ($a(AT_i)$) which aggregate the different attributes. We will assume that bidders do not suffer unexpected incidents ($AT' = AT^t$) and that $a(AT_i) \neq a(AT_i')$ implies a lie in one of the attributes.

Assuming that a bidder utility corresponds to its benefits, its utility function is given by the following expression:

$$u_i(B_i, p) = \begin{cases} p - b_i^t & \text{if } p \neq 0 \\ 0 & \text{otherwise} \end{cases} \quad (3.39)$$

we can obtain the following inequation systems when the product (Equation 3.40), the sum (Equation 3.41) and the weighted sum (Equation 3.42) are used the evaluation function:

- Inequation system when using product as evaluation function

$$\left. \begin{array}{l} \text{(a)} \quad a(AT_1) \neq a(AT'_1) \vee b_1 \neq b'_1 \\ \text{(b)} \quad b_1 * a(AT_1) < b_2 * a(AT_2) \\ \text{(c)} \quad win = \begin{cases} 1 & \text{if } (b_2 * a(AT_2)) > b_1^t * a(AT'_1) \\ 0 & \text{otherwise} \end{cases} \\ \text{(d)} \quad win * \left(\frac{b_2 * a(AT_2)}{a(AT'_1)} - b_1^t \right) < \left(\frac{b_1 * a(AT_1)}{a(AT'_1)} - b_1^t \right) \end{array} \right\} \quad (3.40)$$

where Eq.3.40a defines that bidder 1 is lying about the attributes it bids or its economic price, Eq.3.40b indicates that bidder 1 is the winner of the auction, *win* in Eq.3.40c defines if the bidder would have won the auction if it had told the truth and Eq.3.40d compares the obtained utility and the utility it would have obtained by bidding truthfully.

- Inequation system when using sum as evaluation function:

$$\left. \begin{array}{l} \text{(a)} \quad a(AT_1) \neq a(AT'_1) \vee b_1 \neq b'_1 \\ \text{(b)} \quad b_1 + a(AT_1) < b_2 + a(AT_2) \\ \text{(c)} \quad win = \begin{cases} 1 & \text{if } (b_2 + a(AT_2)) > b_1^t + a(AT'_1) \\ 0 & \text{otherwise} \end{cases} \\ \text{(d)} \quad win(b_2 + a(AT_2) - a(AT'_1) - b_1^t) < (b_1 + a(AT_1) - a(AT'_1) - b_1^t) \end{array} \right\} \quad (3.41)$$

where Eq.3.41a to d follow the same structure as Equation 3.40.

- Inequation system when using weighted sum as evaluation function:

$$\begin{array}{l}
 \text{(a)} \quad a(AT_1) \neq a(AT'_1) \vee b_1 \neq b'_1 \\
 \text{(b)} \quad \mu_1 b_1 + \mu_2 a(AT_1) < \mu_1 b_2 + \mu_2 a(AT_2) \\
 \text{(c)} \quad \text{win} = \begin{cases} 1 & \text{if } (\mu_1 b_2 + \mu_2 a(AT_2)) > \mu_1 b'_1 + \mu_2 a(AT'_1) \\ 0 & \text{otherwise} \end{cases} \\
 \text{(d)} \quad \text{win} \left(\frac{\mu_1 b_2 + \mu_2 a(AT_2) - \mu_2 a(AT'_1)}{\mu_1} - b_1^t \right) < \left(\frac{\mu_1 b_1 + \mu_2 a(AT_1) - \mu_2 a(AT'_1)}{\mu_1} - b_1^t \right) \\
 \text{(e)} \quad \mu_1 + \mu_2 = 1 \\
 \text{(f)} \quad 0 < \mu_1 < 1 \\
 \text{(g)} \quad 0 < \mu_2 < 1
 \end{array} \quad \left. \vphantom{\begin{array}{l} \text{(a)} \\ \text{(b)} \\ \text{(c)} \\ \text{(d)} \\ \text{(e)} \\ \text{(f)} \\ \text{(g)} \end{array}} \right\} \quad (3.42)$$

where Eq.3.42a to d follow the same structure as Equation 3.40 and Eq.3.42e to f delimit the values of μ between 0 and 1.

Satisfiability test

To test the satisfiability of the inequation systems 3.40 to 3.42, , given that they involve both boolean and real arithmetic logic, we need to use a *Satisfiability Modulo Theories* constraint solver [5]. Particularly we used Microsoft Z3 solver [23].

When trying to determine if an inequation system is satisfiable, three results can be obtained: satisfiable, meaning that the inequation system can be solved and that the mechanism is not incentive compatible; unsatisfiable, meaning that the system cannot be solved thus the best strategy is truthful bidding; and unknown. Unknown means that the solver cannot obtain a result in a reasonable time and we cannot determine whether or not the mechanism is incentive compatible. Moreover, the variable domains and the type of logic used can be bounded in order to facilitate the problem solution.

The Z3 program found that all the defined inequation systems are unsatisfiable, pointing that, as stated in the previous subsection, with these three evaluation functions PUMAA remains incentive compatible.

3.5 Considerations Regarding Other PUMAA Properties

The creation of PUMAA is the result of one of the main goals of this thesis: designing an incentive-compatible multi-attribute auction mechanism for workflow resource allocation. How-

ever, as stated in Chapter 2, an auction can be evaluated in terms of different properties. This section analyzes PUMAA's performance regarding such properties:

Efficiency

Taking into account that the winning bid of PUMAA is the one which minimizes $V_0(B_i)$ and, consequently, maximizes the auctioneer utility, PUMAA is an efficient mechanism. Moreover, this minimization ensures that the winning bid is not dominated by any other bid [8].

Buyer optimality

As for any second-price payment mechanism, PUMAA is not buyer optimal because the amount the auctioneer pay may be higher than the amount the bidder requested. Therefore, the auctioneer does not pay the best available price in the market. However, as stated in [49], this is an unavoidable drawback for this type of mechanism to be incentive compatible.

Individual-rationality

This property is strictly related to the participation requirement defined in Section 3.4. Given that the mechanism satisfies the participation requirement, an agent participating in the mechanism does not obtain negative utility simply by bidding in the auction. This means that a bidder which is able to participate in an auction has no reason for not participating in it. Therefore, PUMAA can be considered to be individual-rational.

Budget-balance

In PUMAA the transfer between the auctioneer and the bidders is equal to 0, the winning bidder receives all the money paid by the auctioneer. Thus it can be considered a budget-balanced or ex-post efficient auction mechanism.

Social welfare

On the one hand, following PUMAA, resultant allocations are configured taking into account only the utility of the auctioneer. On the other hand, bidders have no incentive to take into account the utility that is obtained by the other bidders. This clearly demonstrates that PUMAA

follows an utilitarian approach where the resulting allocations maximize the sum of the participants utilities, without taking into account the distribution of such utilities.

Reliability

PUMAA encourages agents to improve their attribute estimations and to be reliable in the task delivery because this will increase their chances of winning the auction, and reduces the chances of them suffering payment reductions. Therefore, bidders obtain higher utilities if they are reliable: if bidders overestimate their attributes they have fewer chances to win the auction whilst underestimating the attributes will imply a payment reduction as the attributes will be delivered worse than agreed. However, the mechanism does not provide any direct reward (better payments or higher chances of winning the auction) to those bidders which have proven their reliability in past auctions.

Robustness

PUMAA can be considered a robust mechanism because, despite the fact it does not ensure that a task will be delivered as it was agreed, the auctioneer is compensated if this occurs. In this way, the utility the auctioneer would have obtained with a proper task execution is preserved. Nevertheless, in future works it would be interesting to improve PUMAA towards the stronger notion of robustness which provides a solution to the bid withdrawal problem [12]. This robust behavior also ensures the quality of service expected for the auctioneer as the bidder is encouraged to keep the agreed quality, otherwise, the bidder would see its utility reduced and the auctioneer would be compensated (preserving its utility).

3.6 Summary

This chapter has presented PUMAA, a reverse multi-attribute auction for allocating resources in workflow environments. PUMAA follows a Vickrey auction schema. As a result it performs its payment once the auctioned task has been performed. Payment is delivered on the basis of the quality of the offered bids: if the winner respects the attributes it promised during the auction it receives a price based on the second best bid. Otherwise, its payments are reduced in order to compensate the difference between the attributes offered and those delivered.

This duality ensures the incentive compatibility of PUMAA whilst, in comparison to other mechanisms, it does not assume that the tasks will be delivered successfully. Moreover, the

payment compensation in case of not providing the bid attributes reduces the harm to the auctioneer due to the task underperformance preserving its utility and saving part of the auctioneer's budget, which then can be used to obtain better suppliers in future tasks.

After describing PUMAA's behavior we have defined the characteristics which a function needs in order to be used as PUMAA's evaluation function (a cornerstone of the mechanism): monotonicity, real-valuation and bijection. The study of these characteristics is performed for PUMAA, however, its results can also be applied in other auction mechanisms like the ones presented in [15] and [65].

Finally the incentive compatibility and other properties regarding PUMAA have been analysed. This study has pointed out that PUMAA is an incentive compatible, efficient, individual rational and budget balanced mechanism which produces utilitarian allocations.

A FRAMEWORK FOR MULTI-ATTRIBUTE AUCTION CUSTOMIZATION

Mechanism designers face the challenge of maximizing participants' utilities but they are also required to fulfill other requirements depending on the domain they are dealing with (e.g. incentive compatibility, robustness, efficiency, etc.). However, as discussed in Chapter 2 some of these properties are in conflict due to their contradictory nature.

PUMAA offers a wide variety of uses whilst guaranteeing a set of properties (incentive compatibility and efficiency). However, depending upon the domain, the mechanism designer may be interested in enhancing other properties even if it means sacrificing some of the existing ones (e.g. an auctioneer may be more interested in obtaining robust and reliable allocations than in obtaining efficient ones [65, 69]). With this in mind, in this chapter we propose a PUMAA-based framework for multi-attribute auction customization: FMAAC (framework for multi-attribute auction customization). In this way, the auction mechanism presented in the previous chapter is parameterized in order to provide a general framework, which in turn can be used to reach domain-dependent requirements.

For this purpose we first analyze the kind of attributes which can take part into an auction mechanism and the role they play in the auction. With this information we then build the general framework FMAAC.

4.1 Attribute Typologies in Multi-attribute Auctions

In multi-attribute auctions, the item which is sold is defined due to a set of attributes, in addition to the price, that it is determined in the auction process. Multi-attribute auctions may

involve many attributes regarding the item, as for example, the size and quality requested by the buyers, and the price offered by the sellers.

Moreover, in some uni-attribute auctions the auction is cleared by taking into account other attributes besides the ones offered by the bidder. For instance, information about bidders' past performance [69] (reliable or not), previous auction results [54] (the number of auctions which an agent has won) or opinions which the auctioneer has in regard to bidders [84] can be taken into account when deciding the forthcoming winners.

In this section we consider all of the above attributes, as they all influence the winner determination problem and the payment mechanism of multi-attribute auctions. In doing so we distinguish two main criteria: attribute *ownership*, and *verifiability*. Ownership means that attributes can be characterized according to the information source (the auctioneer itself or the participating bidders). On the other hand verifiability concerns the capability of an attribute to be verified by an agent.

4.1.1 Ownership of the attributes

The ownership of the attributes refers to the source of the attributes, namely, which agent is introducing a given attribute to the auction mechanism. Thus, we can classify attributes as auctioneer-provided or bidder-provided:

- **Bidder-provided attributes:** This is the set of attributes that bidders express regarding the qualification of the goods to be served. These attributes are the ones which define bids and are used in every kind of auction: uni-attribute auctions (1 bidder-provided attribute) and multi-attribute (n bidder-provided attributes for each bid). For instance, these attributes are the ones used in the classical Vickrey auctions or in PUMAA.
- **Auctioneer-provided attributes:** This is the set of attributes which the auctioneer uses to extend the bidder's bid. They can be used to qualify bidders or their bids on behalf of a given objective criteria. These attributes can involve: number of withdraws on behalf of a given winner, number of times the bidder has won an auction, number of times the bidder reported a lower quality than the one it bid for, etc. Moreover, these attributes can be used to express the subjective auctioneer beliefs regarding bidders or their bids. For instance, they can be used to define reputations or how popular a bidder is according to the auctioneer's perspective. Auctioneer-provided attributes may appear in both uni-attribute [84] and multi-attribute [69] auctions. An example of these attributes can be found in the Google Sponsored Search Auctions (see Chapter 2) where the auctioneer

	Verifiable	Unverifiable
Bidder Ownership	<p>Verifiable bidder provided attributes</p> <p>Delivery times, qualities, energy consumptions, CO₂ emissions...</p>	<p>Unverifiable bidder provided attributes</p> <p>Economic cost, CO₂ emission quota</p>
Auctioneer Ownership	<p>Auctioneer provided attributes</p>	
	Auctions won, past performance ...	Agent's reputation

Figure 4.1: Classification of the attributes involved in a multi-attribute auction according to its origin and verifiability.

adds the quality of past advertisements to bids. These attributes can be used to model the auction behavior (e.g. to obtain egalitarian social welfares, or to build reliable auctions).

Since it can be assumed that an auctioneer will not try to deceive itself, we can say that auctioneer-provided attributes are trustworthy. Thus the auctioneer does not need to concern itself with the reliability of these attributes.

4.1.2 Verifiability

The verifiability of an attribute concerns the capacity of an agent to check the truthfulness of the attribute by means of an objective measure. Thus an attribute can be verifiable or unverifiable.

- **Unverifiable attributes:** are the set of attributes defined by an agent whose true values are only known by the agent itself. These attributes are also the ones which define the auction currency. A typical example of this kind of attribute is the economic value which a bidder offers or asks to obtain an item or for providing a service. The bidder

knows its true value, however the auctioneer has no way to know the true value of the attribute neither before nor after delivering the auctioned item. Similarly, bidders could trade permits for generating certain amounts of CO_2 emissions (e.g. as happens with emissions trading between countries) [36].

Despite the fact that there can be more than one unverifiable attribute, auctions are typically designed using only one unverifiable attribute due to the complexity of introducing more than one¹. In cases where there is more than one currency attribute [68] (e.g. international auctions), all the attributes are translated into a unique attribute or currency so the auctioneer only deals with one type of unverifiable attribute (for example a currency which acts as a standard monetary unit or a virtual currency). Therefore, despite the existence of more than one unverifiable attribute during the bidding process, the auction is performed considering just one single unverifiable attribute. From this point of view, when we refer to an unverifiable attribute we mean the unique attribute which acts as currency for the auction (namely the attribute itself or the virtual currency which aggregates more than one unverifiable attributes).

Unverifiable attributes appear in all the auction types. In uni-attribute auctions they correspond to the entire bid offered by bidders, whilst in multi-attribute auctions they correspond to at least one of the attributes provided in the bid. For example, in PUMAA, the economic amount is the unverifiable attribute.

- **Verifiable bidder attributes:** These are the set of attributes which are defined by agents whose true value can be known and checked by another agent. It does not matter if they can be verified before or after delivering the auctioned item, however, this verification must always be completed before the payment is performed. Examples of this type of attribute are delivery times, electricity consumption or other physical specifications. For example when auctioning a task which will be carried out by a bidder, the bidder can specify a certain delivery time t . Once the task has been completed the auctioneer can then check if the final delivery time t' was as specified during the bidding process. As these attributes can be checked, they can be used to adjust payment to bidders or to establish parameters to describe bidder qualities (auctioneer-provided attributes).

Since uni-attribute auctions are composed only by unverifiable attributes, these kinds of attribute only appear in multi-attribute auctions.

¹observe that Section 3.3.1 specifies that the evaluation function of an score-based auction must be bijective regarding the currency attribute.

Attribute Type	Unverifiable bidder-provided	Verifiable bidder-provided	Auctioneer-provided
Vickrey auction [85]	Economic cost		
Google PPC auction [84]	Pay per click		Ad Quality
Priority auctions ² [54]	Economic cost		Priority
Che's second-score a. [15]	Economic cost	Attribute bundle	
Porter's f.t.a. [65]	Economic cost, POS		
PUMAA	Economic cost	Delivery time, energy, etc.	

Table 4.1: Attributes involved in some of the auctions presented in this thesis.

As mentioned in the previous section, taking into account that auctioneer agents are in charge of determining the auction winners, it can be assumed that they will not try to deceive themselves. Consequently, the auctioneer-provided attributes do not need to be considered in this second classification. Figure 4.1 shows, that attributes involved in auctions can be classified as *auctioneer-provided attributes*, *verifiable bidder-provided attributes* and *unverifiable bidder-provided attributes*. Table 4.1 shows the types of attributes involved in some of the auctions mentioned along this thesis.

In Figure 4.2 we illustrate a simple multi-attribute auction in which the three different types of attributes are used. Auctioneer A calls an auction in order to find an agent to perform a certain task. Bidders send bids containing the economic cost, they expect to charge for the task (attribute b) and the delivery time they propose (dt). Once the bidders have sent their proposals, A will include a reliability attribute r which rates its satisfaction regarding previous deals with the different bidders. Finally, the winner of the auction is computed using every attribute involved in the auction process (b, dt, r). In this example b is an unverifiable bidder-provided attribute as the auctioneer cannot know the true value of the cost attribute for each bid, at most, it can estimate it. dt is a verifiable bidder attribute as, once the task is finished, the auctioneer can compare the real delivery time dt' with the one which was provided in the bid. Finally, r is an auctioneer-provided attribute as it is added to the bid by the auctioneer itself and it is also used to determine the winner of the auction.

²This method is reviewed in Chapter 5




Figure 4.2: Example of the attribute types used in the bidding process: Cost is a unverifiable bidder provider attribute, delivery time a verifiable bidder-provided attribute and reliability is an auctioneer-provided attribute.

4.2 FMAAC: Framework for Multi-attribute Customization

Using the classification of attributes provided in the previous section we propose to generalize PUMAA by incorporating auctioneer-provided attributes. This addition allows designers to cover a wide range of problems which PUMAA is not prepared to deal with. For instance, the use of auctioneer-provided attributes allows to obtain egalitarian allocations instead of utilitarian ones. Another possible extension can consist of using the auctioneer-provided-attributes to favor those agents which have provided high performance in past auctions. Thus obtaining more reliable allocations. Depending upon the kind of auctioneer-provided attributes employed, the auction mechanism can be customized towards different goals or properties.

Given that different types of auctioneer-provided attributes results in different types of auctions with different properties, we can no longer define this as an auction mechanism but as an auction framework. We call this framework FMAAC: Framework for Multi-Attribute Auction Customization.

FMAAC takes PUMAA and introduces three main modifications. The first and most obvious one is the inclusion of the auctioneer-provided attributes (from now on A^P) which concern the WDP and the payment mechanism that must be adapted in order to include A^P . The second is

the differentiation between verifiable bidder-provided attributes (A^v) and unverifiable bidder-provided attributes (A^u) which affects the whole auction protocol. Finally, it includes a new *Attribute Information Update* step in the auction protocol in order to compute the A^p values.

To describe FMAAC we assume that auctions succeed over time so auctioneers can obtain information about the participant agents. As in the case of PUMAA, we keep the assumption that there are no externalities in the process, as well as no budget constraints by any agent and that bids are presented under sealed bid.

Regarding the attribute typologies described in the previous section, we adopt the following notation:

- Auctioneer provided attributes of the agent a_i :

$$A_i^p = (at_{1i}^p, \dots, at_{m_p i}^p)$$
 with $at_{1i}^p \in D_1^p, \dots, at_{m_p i}^p \in D_{m_p}^p$
- Verifiable (bidder-provided) attributes of the agent a_i :

$$A_i^v = (at_{1i}^v, \dots, at_{m_v i}^v)$$
 with $at_{1i}^v \in D_1^v, \dots, at_{m_v i}^v \in D_{m_v}^v$
- Unverifiable (bidder-provided) attributes of the agent a_i :

$$A_i^u = (at_i^u)$$
 with $at_i^u \in D_1^u$

where D_j^x is the domain of an attribute a_j^x .

4.2.1 Call for proposal - Defining the bidder-provided attributes

When an agent wants to buy or sell an item, it summons an auction. An item is not necessarily a physical object. For example, it can be a service externalization, or the need of a resource to perform a task. The auctioneer desiring to buy or sell a task will define if the auction is reverse or not. As this thesis is specially focused on resource and task allocation, we will refer only to reverse auctions. However, FMAAC can be utilized in both reverse and forward auctions.

The auctioneer tries to obtain the best option at the best possible price. For that purpose, the agent builds a call for proposal *CFP* defining the item to be purchased to the bidders. In the *CFP* the auctioneer will define the item *it* it wants to buy but also the attributes AR_0 related to *it* it wants to buy but also the attributes AR_0 related to *it* in which it is interested. AR_0 will define the attributes which each bidder will have to include in its bid. In other words, AR_0 stipulates the bidder-provided attributes involved in the auction (both A^v and A^u). Thus, in the first step of the auction, the auctioneer defines the verifiable attribute which will act as the auction currency (A_0^u), and the set of verifiable attributes A_0^v which will also be taken into account during the auction process.

$$CFP = (it, AR_0)$$

$$CFP = (it, A_0^u, A_0^v) \quad (4.1)$$

The auctioneer then sends this CFP to all the agents within the market and announces it will accept bids during a certain time frame.

4.2.2 Bidding - The role of unverifiable attributes

In the bidding stage, agents which own items to be sold or rented which match the description of it must decide if they are interested in the auctioneer's proposal and which bids they will submit.

When an agent a_i receives a call for proposals CFP it assesses whether it is interested in participating in the auction. For that purpose it determines its own capacity to attend to the request $(it, A_0^v \subseteq D_1^v \times D_{m_v}^v)$, with $A_i^v = (at_{1i}^v, \dots, at_{m_v i}^v)$. As the auction follows a sealed bid schema and the bidders do not know what bids its contenders offered in past auctions, the bidder is only concerned with its own offer. In cases where the bidder is interested in participating, it needs to return a bid it considers it will maximize its expected utility. It also includes the verifiable attributes it can provide but also the unverifiable attributes (usually the price it requires) which the auctioneer indicated.

The true value of the unverifiable attribute is conditioned by the item for sale and the set of verifiable attributes. For instance, for a given courier company, the cost of delivering a package can be defined by the salary of the employee doing the delivery, the fuel required for the transportation, a proportion of the company's fixed costs and the type of transport which will use (e.g. a faster transport may imply higher costs). Thus, the valuation which a bidder a_i makes of the unverifiable attributes depends on the item to be sold and its verifiable attributes:

$$A_i^u = v_i(it, A_i^v) \quad (4.2)$$

Therefore, the bid can be defined as follow:

$$B_i = A_i^u \oplus A_i^v = (at_i^u, at_{1i}^v, \dots, at_{m_v i}^v) \quad (4.3)$$

Cheating agents provide bids with values other than their true valuation. For example, with a bid $B_{i'} = (at_{1i'}^u, at_{1i'}^v, \dots, at_{m_v i'}^v)$ where either $A_i^v \neq A_{i'}^v$, $A_{i'}^u \neq v_i(it, A_i^v)$ or both, the agent could

be offering better attributes than its skills or lower price². This has the aim of deceiving the mechanism and winning the auction with better economic conditions.

4.2.3 Winner determination - The role of auctioneer-provided attributes

Upon the receipt of the bids, the auctioneer has the ability to extend the received bids with auctioneer-provided attributes A^P . A^P can include different kinds of information and opinions regarding the bidders based on past auctions. Thus, the auctioneer extends each received bid with the information it has recorded in past, obtaining a modified bid B' which contains all the attributes which will be used to determine the auction winner and its payment:

$$B'_i = B_i \oplus A_i^P \quad (4.4)$$

$$B'_i = (at_i^u, at_{1i}^v, \dots, at_{m_v i}^v, at_{1i}^p, \dots, at_{m_p i}^p) \quad (4.5)$$

where A_i^P are the auctioneer-provided attributes, $at_{1i}^p, \dots, at_{m_p i}^p$.

Once the different bids have been extended, the auctioneer needs to determine the winning bid, maximizing its expected utility. The expected utility must be coherent with the auctioneer's utility function³. It is defined as the difference between the valuation v_0 the auctioneer gives to the item it to be bought and the result of an aggregation function V_0 which evaluates all the attributes of the auction (including the auction currency and the auctioneer-provided attributes):

$$\bar{u}_o(it, B'_i) = v_0(it) - V_0(B'_i) \quad (4.6)$$

$$\bar{u}_o(it, AT_i^u, AT_i^v, AT_i^p) = v_0(it) - V_0(AT_i^u \oplus AT_i^v \oplus AT_i^p) \quad (4.7)$$

where $v_0(it)$ is a function which describes the value which a_0 gives to the item it it wishes to buy and $V_0(AT_i^v \oplus AT_i^u \oplus AT_i^p)$ the value which it gives to the set of attributes bidded.

To evaluate the bids received the auctioneer uses an aggregation function V_0 which must be monotonic, real-valued and bijective (see Chapter 3). Therefore, the winner of the auction is the bid which minimizes the value of V_0 :

$$winner = argmin_i(V_0(AT_i^u \oplus AT_i^v \oplus AT_i^p)) \quad (4.8)$$

If we assume that there is only one bidder-provided attribute which acts as the auction currency (see Sections 4.1.2 and 4.2.2) we can define at_i^u as the auction currency ($at_i^u = b_i$)

²As with PUMAA, we assume that bidders are able to accurately estimate their skills

³We consider the same auctioneer's utility function as in the previous chapter, see Equation 3.12

and express the winner determination problem as:

$$\operatorname{argmin}_i \left(V_0 \left(b_i, at_{1i}^v, \dots, at_{m_v i}^v, at_{1i}^p, \dots, at_{m_p i}^p \right) \right) \quad (4.9)$$

Logically, the number of parameters which the aggregation function V_0 accepts must correspond to the number of attributes involved in the auction ($|AT^u| + |AT_i^v| + |AT^p|$).

Depending upon the type of auctioneer-provided attributes A_i^p incorporated in the auction, these attributes may be used to represent auctioneer opinions (such as trust in a bidder based upon past performances) or information regarding the bidders (for instance, the number of times an agent has won or participated in past auctions). According to Equation 4.8, the attributes used to extend the bids will affect the characteristics of the auction and the resulting allocations. For example, below we present two examples of FMAAC-based mechanisms which can be used to obtain specific types of allocations:

- **fair-PUMAA - Egalitarian social welfare:** A priority attribute which defines bidder's history (number of auctions won, lost and participated) is used in order to enhance the equity and fairness of the allocations. This type of allocation can be useful to keep bidders interested in future auctions and to avoid recurrent auction problems such as the bidders drop problem [54, 47].

For instance, the auctioneer can assign a priority attribute $w_i \in [0, 1]$ to each bidder, according to the ratio between the number of auctions in which they have participated, and the auctions they have lost within a specific time window. The higher the ratio of lost auctions, the higher the priority of the bidder, meaning that a bidder with high priority should be awarded a task soon to prevent the agent from leaving the market due to too low income and obtaining an egalitarian social welfare.

- **trust-PUMAA - Reliability:** A trust attribute $\tau_i \in (0, 1]$ defines the performance of each bidder in past auctions in order to improve the reliability of the allocations. τ_i may express the relation between the number of auctions won and the number of successful tasks (or successfully delivered items), in this way unreliable bidders would have a lower chance of winning future auctions.

It is important to take into account that the use of auctioneer-provided attributes may affect bidders' behavior. For instance, the use of a trust attribute can incentivize bidders to improve their accuracy when estimating their delivery times but an abuse of this attribute may lead certain bidders to abandon the auction due an unexpected estimation error at an early stage

of the allocation process. In the same way, whilst the use of priorities can preserve the number of bidders within an auction market, an abuse priorities may lead agents to provide dummy bids [72] (low bids submitted with the purpose of losing an auction) in order to increase their probability of winning a future auction.

Moreover, PUMAA properties such as efficiency or incentive compatibility can also be affected by the use of A^P . For instance, the use of a priority attribute may result in the bidder with the best combination of price (A^V) and verifiable attributes (A^U) losing the auction due to a low priority. This breaks down the efficiency property. In this case, efficiency is sacrificed for the sake of egalitarian social welfare or long term efficiency (maximizing the utility in the long run, not in a single-shot auction). Thus, we recommend that mechanism designers study of what bearing a given auctioneer-provided attribute might imply before adding it to the auction mechanism.

4.2.4 Payment - Playing all the attributes together

The payment mechanism of FMAAC is computed after the auctioned item is delivered. Using verifiable attributes, the auctioneer checks if the item has been delivered under the same or better (\succeq) conditions as agreed or in worse ones (\prec). In the first case, the auctioneer pays the amount p which the winning bidder should have bid in order to obtain the same score V_0 (taking into account also the auctioneer-provided attributes) as the second best bid:

$$V_0(p_1, at_{1,1}^v, \dots, at_{m_v,1}^v, at_{1,1}^p, \dots, at_{m_p,1}^p) = V_0(b_2, at_{1,2}^v, \dots, at_{m_v,2}^v, at_{1,2}^p, \dots, at_{m_p,2}^p) \quad (4.10)$$

which we simplify to:

$$V_0(p_1, AT_1^v, AT_1^p) = V_0(b_2, AT_2^v, AT_2^p) \quad (4.11)$$

In the second case, the bidder will receive the amount p it should have bid together with the delivered verifiable attributes $AT_i^{v'} = at_{1,i}^{v'}, \dots, at_{m_v,i}^{v'}$ in order to match the valuation of the bid it originally offered:

$$V_0(p_1, at_{1,1}^{v'}, \dots, at_{m_v,1}^{v'}, at_{1,1}^p, \dots, at_{m_p,1}^p) = V_0(b_1, at_{1,1}^v, \dots, at_{m_v,1}^v, at_{1,1}^p, \dots, at_{m_p,1}^p) \quad (4.12)$$

which we simplify to:

$$V_0(p_1, AT_1^{v'}, AT_1^u) = V_0(b_1, AT_1^v, AT_1^u) \quad (4.13)$$

Thus, since the aggregation function V_0 returns a value x with the set of attributes b , AT^v and AT^p ($V_0(b, AT^v, AT^p) = x$), if we define $V_0^{-1}(x, AT^v, AT^p)$ as the inverse function which given a value x returns the b value of V_0 , the payment can be defined as:

$$p_1 = \begin{cases} V_0^{-1}(V_0(b_2, A_2^v, A_2^p), A_1^v, A_1^p) & \text{if } A_1^v \geq A_1^p \\ V_0^{-1}(V_0(b_1, A_1^v, A_1^p), A_1^v, A_1^p) & \text{if } A_1^v < A_1^p \end{cases} \quad (4.14)$$

4.2.5 Attribute information update

At this stage, the auctioneer updates its auctioneer-provided attributes based on objective information. This information can concern a wide range of domains; therefore, it may need to collect information at a different stages during the auction depending on the type of attribute being handled.

For instance, when using priorities based on the results of previous auctions, the auctioneer updates their values as new information appears regarding the number of auctions which each participant has won and lost. On the other hand, when using a trust value, after receiving the auctioned task the auctioneer can update the attribute which defines the trust value of the winning agent according to how the task has been performed.

4.2.6 A simple example

To illustrate the behavior of FMAAC we present an example of the instantiation of FMAAC in order to increase the reliability of the auction allocations by PUMAA (trust-PUMAA [80]).

In this case an auctioneer a_0 needs to externalize a service S_0 which must be finished before a certain time t_0 . The earlier S_0 is finished, the higher the auctioneer satisfaction or utility will be, as it will have more time to invest in the rest of the tasks.

We analyze the first case following the PUMAA approach with verifiable and unverifiable attributes. In this case, the utility of S_0 is defined on the basis of the benefits the auctioneer obtains from having S_0 finished ($v_0(S_0) = 50\text{€}$). The auctioneer wants S_0 to be finished before 100 minutes have elapsed ($t_0=100$), moreover, it values each minute which it can save in S_0 as 2€ . Thus, given the economic price p , the real delivery time t'_i , the economic cost b_i and delivery time t_i offered in a bid, the auctioneer's utility and expected utility can be defined as follows:

Bidder	B_i	A^u			A^v		A^p		
		b_i	t_i	τ_i	s_i	s_i^{ok}	$B'_i (B_i \oplus \tau_i)$	$V_0(B'_i)$	rank
a_a	(30, 92)	30	92	0.60	10	6	(30, 92, 0.60)	$23.\hat{3}$	2
a_b	(40, 90)	40	90	0.80	5	4	(40, 90, 0.80)	25	3
a_c	(30, 95)	30	95	0.90	10	9	(30, 95, 0.90)	$22.\hat{2}$	1

Table 4.2: List of bids and their corresponding ranks and evaluations.

$$u_0(S_0, p_i, t'_i) = v_0(S_0) - f_0(p_i, t'_i) \quad (4.15)$$

$$\bar{u}_0(S_0, b_i, t_i) = v_0(S_0) - f_0(b_i, t_i) \quad (4.16)$$

The evaluation function V_0 could be defined as:

$$V_0(b_i, t_i) = b_i - (100 - t_i) * 2 \quad (4.17)$$

It can be observed that b_i is an unverifiable bidder-provided attribute referring to the economic cost of performing S_0 and where the delivery time t_i is a verifiable attribute.

Consider now that the auctioneer must accomplish other tasks which depend on the services it externalizes, and it is interested its outsourced services being delivered on time. Thus, it decides to favor those agents which have provided tasks without delays in the past. For this purpose we need to consider an auctioneer-provided attribute as trust, something that PUMAA does not allow. Conversely, with FMAAC, the auctioneer can use a trust attribute τ_i which corresponds to the relationship between the number of services s_i an agent a_i has performed with the number of services s_i^{ok} it has delivered on time:

$$\tau_i = \frac{s_i^{ok}}{s_i} \quad (4.18)$$

Therefore, the auctioneer also needs to take into account the auctioneer-provided attribute, trust, in the evaluation function V_0 . An example of this addition could be the follow:

$$V_0(b_i, t_i, \tau_i) = (b_i - (100 - t_i) * 2) * \frac{1}{\tau_i} \quad (4.19)$$

Consider that the auctioneer receives three different bids B_i from bidders a_a , a_b and a_c exposed in Table 4.2.

After receiving the bids, the auctioneer ranks them from the lowest to the highest score, and determines that the auction winner is a_c . Note that without a trust attribute, the winner would

have been a_a as it has a lower delivery time; however, in this case, as we are also considering the trust attribute τ , a_c obtains a lowest rating.

Finally, when the a_c finishes S_0 , the auctioneer evaluates the obtained result and proceeds to pay the bidder. As the payment mechanism considers two situations, we illustrate two different payment scenarios: when the delivered items meets the auction agreement and when the bidder breaks it.

- If a_c finishes S_0 with a delivery time equal or lower than the offered (t_c), it will receive the amount it should have offered to equal the second best bid (bider a_a). So the payment would be:

$$V_0(p_c, t_c, \tau_c) = V_0(b_a, t_a, \tau_a) \quad (4.20)$$

$$p_c = 23.\hat{3} * 0.9 + (100 - 95) * 2 = 31 \quad (4.21)$$

After the payment process the auctioneer would have to update its information regarding the auctioneer-provided attributes. Given that the task had been successfully completed, the auctioneer updates the information regarding the number of successful tasks which the winning bidder has performed ($s_c = 11$ and $s_c^{ok} = 10$). Thus, the agent a_c 's trust increases to $\tau = 0.909$, improving the chances of a_c of winning further auctions.

- Otherwise, if a_c finishes S_0 later than bidded ($t'_c < t_c$)⁴, it will receive the amount it should have bid to achieve the same evaluation it obtained with the original bid. To illustrate that, we will consider $t'_c = 100$.

$$V_0(p_c, t'_c, \tau_c) = V_0(b_c, t_c, \tau_c) \quad (4.22)$$

$$p_c = 22.\hat{2} * 0.9 + (100 - 100) * 2 = 20 \quad (4.23)$$

In this situation, the bidder payment decreases due to the failure of the bidder to complete the task within the agreed time. After the payment process, the auctioneer updates its information regarding the number of successful tasks which the winning bidder has performed ($s_c = 11$ and $s_c^{ok} = 9$). Thus, the agent a_c 's trust is reduced to $\tau = 0.818$, reducing the chances of a_c of winning further auctions.

⁴notice that $<$ means *worse than*

4.3 Summary

This chapter has discussed the fourth and fifth contributions of this PhD thesis. First, we analyzed the kind of attributes which are involved in auctions. These are unverifiable bidder-provided attributes which are those which act as the auction currency and whose true values is only known by the bidder which offered them. Verifiable bidder-provided attributes are the set of attributes offered by bidders whose true value can be verified after the item is delivered. Auctioneer-provided attributes are introduced into the auction by the auctioneer itself and can be used to express information and beliefs of the auctioneer regarding the bidders.

Using this classification, we presented FMAAC, a framework based in PUMAA for customizing multi-attribute auctions. This framework takes the idea of adding auctioneer-provided attributes in multi-attribute auctions from the Google's Ad auctions (a uni-attribute auction). In FMAAC the auction designer can add a set of attributes for defining the auctioneers' opinions regarding the bidders in the market. The use of these attributes, which are utilized during the entire auction protocol, endows the auction with new properties such as egalitarian social welfare, reliability or any other property desired by the mechanism designer. However mechanism designers should add auctioneer-provided attributes with caution as they can act as a mixed blessing. The addition of these attributes may condition bidding strategies and may affect other auction properties such as incentive compatibility or efficiency. Thus, we recommend designers assess the implications of adding a new attributes to the auction in detail before proceeding with their incorporation

MULTI-DIMENSIONAL FAIRNESS FOR MULTI-ATTRIBUTE RESOURCE ALLOCATION

The auction designer's goals include optimizing the payoff or revenue of bidders and auctioneers, so that all the participating agents are satisfied and remain in the market place. To evaluate how satisfied the bidders are with the auction outcome, as well as the revenue obtained by the auctioneer [16, 26] social welfare measures can be defined. The utilitarian view of social welfare has been the main approach of auctions and consists of aggregating all the agents' outcomes, towards maximizing their payoff or revenue. In this utilitarian approach, such aggregation does not consider the fact that there could be big differences among agents' payoffs. When auctions are repeated over time (recurrent auctions), this situation may lead to the dissatisfaction of certain participants who may eventually decide to leave the market. When this occurs, only the most powerful bidders remain in the market, gaining the chance to create an oligopoly, to control the market price and to provoke a general fall of prices which may bankrupt the auctioneers. Literature often refers to these situation as the bidder drop problem [43] and the asymmetric balance of negotiation power [56].

To tackle this problem, fairness measures have been used in uni-attribute auction design, in what are known as egalitarian social welfare approaches. In this scenario, the behavior of bidders can be totally selfish, as the auctioneer is the only agent concerned with obtaining an egalitarian social welfare. Thus, the auctioneer agent uses fairness measures to distribute the revenues to keep bidders interested in participating. By doing this, the auctioneer sacrifices instantaneous optimality in order to maximize its utility in the long run. This provides

opportunities for disadvantaged bidders (that without fairness might not have the chance of winning any auction) to prevent the most powerful agents from dominating and controlling the market [54]. In the end, more bidders mean a higher level of competition, leading to a more competitive market and allowing disadvantaged bidders to win some auctions is a price to pay in order to avoid the creation of undesired oligopolies.

In egalitarian uni-attribute auctions, fairness has been considered exclusively from an economic point of view (which is the attribute maximized in the auction). Consistently, in multi-attribute auctions, fairness cannot be limited to payoffs and revenues obtained by agents as focusing the application of fairness in a single attribute auction may involve undesirable consequences regarding the rest of the attributes (e.g. unbalance workloads or produce delays). To prevent this issue, we propose the application of fairness mechanisms considering not only the economic aspects of the auction but also the remaining attributes involved in the resource allocation decision making process. We call this multi-dimensional fairness.

In this chapter, we propose the use of FMAAC with a multi-dimensional fairness mechanism based on various priorities in order to increase the social welfare resulting from a large sequence of auction allocations. Particularly, we present a collection of different priority methods for implementing multi-dimensional fairness to a multi-attribute auction mechanism (two qualitative and two quantitative approaches with a deterministic version and a stochastic version of each). Priorities are computed using information regarding all the attributes involved in the resource allocation process, avoiding unwanted behaviors. For example, the situation in which the auctioneer achieves a cheap price but a large delay in performing some tasks (a behavior which other fairness mechanisms based exclusively on price could exhibit).

fair-PUMAA is based upon PUMAA, thus, it uses an aggregation function as an evaluation function and it preserves the utility of the auctioneer in case of a task being underperformed. However, since its principal characteristic is that it favors egalitarian allocations, it can be considered an egalitarian auction (Figure 5.1).

This chapter first introduces basic concepts regarding fairness in auctions. We then present fair-PUMAA, a PUMAA extension built with FMAAC which includes a priority-based multi-dimensional fairness mechanism. Finally, in Section 5.4 we present eight different methods for computing the priorities in fair-PUMAA.

5.1 Preliminary Information

This section introduces some previous literature regarding fairness in auctions which have been taken as starting point to build a multi-dimensional fairness approach.

Despite it has been proven that preserving the number of participants in an auction increases its efficiency in the mid to long term [47, 54] little research has been done in the field of fairness in recurrent auctions. There are two main approaches in tackling this problem: discriminatory price recurring auctions (DPORA) and Priority auctions (both have been applied in uni-attribute auctions). These two methodologies estimate the probability of bidders leaving the auction market; then, they use that probability to favor the weakest agents in order to encourage them to remain in the market. Using this probability, the first approach suggests establishing a reservation price and distributing unsold goods among the agents with highest priorities. The second mechanism proposes using priorities to help the weaker agents win some of the auctions.

Both approaches use a bidder-provided unverifiable attribute (price) and an auctioneer-provided attribute (a priority or an index) in the WDP but neither use attributes other than price in the payment. Therefore, our approach is a novel contribution for multi-attribute auctions.

5.1.1 Discriminatory price optimal recurring auction

Lee and Szymanski [43] propose a fair mechanism for allocating multiple units of the same item in recurrent auctions (MURA). They improve the welfare of the weakest agents by establishing a reservation price so that some goods cannot be sold, if the prices offered by bidders are under this reservation price. The remaining goods are then distributed amongst the agents which are in highest risk of leaving the market.

The work of Lee and Szymanski is based upon the supply and demand principle of microeconomics. The mechanism sets a reservation price based upon the demand obtained at the last instance this kind of item was auctioned. Following this, the bidders offering a higher bid than the reservation price obtain one of the auctioned items, the remaining goods are then distributed among the losing bidders according to the VLLF-BDC algorithm (Valuable Last Loser First Bidder Drop Control). The algorithm classifies the bids according to whether the bidders have increased their bids since the last time the item auctioned was sold (marked bids) or if they have not (unmarked bids). The auctioneer, then, allocates the remaining items to the bidders which have the highest marked bids. If after this allocation there are still goods to be

allocated, the goods are assigned to the highest unmarked bids. The payment then follows a first-price policy (each bidder pays exactly what it bid) but this is not incentive compatible.

Lee [44] proposes a modification by including a Vickrey based payment mechanism (PI-ORA). The author claims that this modification rewards bidders' participation, making this an incentive compatible mechanism.

5.1.2 Priority auctions

Murillo et al. [55] propose a first price, fair method based on priorities. In this article, the auctioneers calculate that the probability of an agent leaving the market due to a bad auction record. A priority is then assigned to each agent according to this probability. Bidders with a poor auction win record have a higher priority, meaning that they have a higher chance of winning the next auction. The authors propose using the priority attribute to condition the auction clearance by aggregating the priority to the bid (e.g. using the product), however, in the payment rule, the priority attribute is omitted and does not alter the payment the winners receive

Although the authors prove that their approach solves the bidder drop problem without incurring a waste of resources (leaving certain goods without a buyer), this mechanism suffers from an asymmetric balance of negotiation power [43, 54]. This is caused by the lack of incentive compatibility even in a one-shot auction (observable when studying an isolated auction from the allocation sequence without taking into account previous or further auctions).

The lack of incentive compatibility in this mechanism is due to the fact that it follows a first price philosophy. The transformation of this approach to a second price mechanism may solve the strategy proofness issue, however, it poses the problem of how the priority should be integrated in the payment system as ignoring priority in a second price system could induce bidders to pay more than what they bid. In our work we follow a similar strategy to the one employed by Murillo et al. We also use priorities to favor the weakest agents in the system, however, we use a second-price auction instead of a first-price one. This ensures incentive compatibility, at least, for one-shot auctions. The Vickrey-based structure of FMAAC allows us to follow a second price structure, as the priority attribute can be taken into account during the whole auction protocol. Moreover, given the dimensionality of the problem we are dealing with, priority is computed with not only the price but also the rest of attributes within the auction taken into account.

5.2 Multidimensional Fairness

Our goal is to define a fairness mechanism that considers the different attributes, contributing to the decision-making, in an auction-based resource allocation process. In particular, we suppose an environment where a set of resource agents compete for tasks that are repeated arbitrarily over time (i.e the supply chain environment). Thus, the auctioneer can keep track of the historical outcomes of each bidder in past auctions, identifying resource agents at risk of leaving the market. Thus, fair decisions can be made to keep agents interested in the market.

To achieve our goal, we propose the assignment of priorities to resource agents according to their historical results in the resource allocation process. Priorities are defined in $w_i \in [0, 1]$. The higher the ratio of lost auctions, the higher the priority of the bidder (due to its high probability of leaving the market). Consistently, a high priority increases the chances of a bidder winning one of the forthcoming auctions.

Priorities can be handled as a bidder attribute in the auction model, in a similar fashion to the click likelihood in sponsored search position auctions. However, position auctions only consider price to characterize tasks whereas we are dealing with several task attributes (price, time to deliver, etc.). Therefore, we use the FMAAC model allowing us to combine both task attributes and bidders attributes.

In a nutshell, when a task needs to be fulfilled an auctioneer calls for an auction and all the interested bidders submit their bids B_i . Bids are vectors composed by all the attributes characterizing the tasks and the price. Next, the auctioneer adds the priority attribute corresponding to each bidder into their bids ($B'_i = B_i \oplus \langle w_i \rangle$). Like any other attribute, priority w_i will bias the winner of the auction and its revenue inasmuch as w_i is taken into account both in the WDP and in the payment mechanism. Once the auction has finished, the auctioneer updates the priorities of the participants according to the results of the auction.

Fairness at the global level, then, is not a matter of a single attribute like price or priority. It is multi-dimensional as it arises from all the attributes involved in the auction: ones characterizing the task (price, quality, delivery times, etc.) as well as the bidders (priority).

5.3 fair-PUMAA

fair-PUMAA is the fair version of PUMAA, which is extended using FMAAC with the aim of incorporating a multi-dimensional fairness mechanism to multi-attribute auctions.

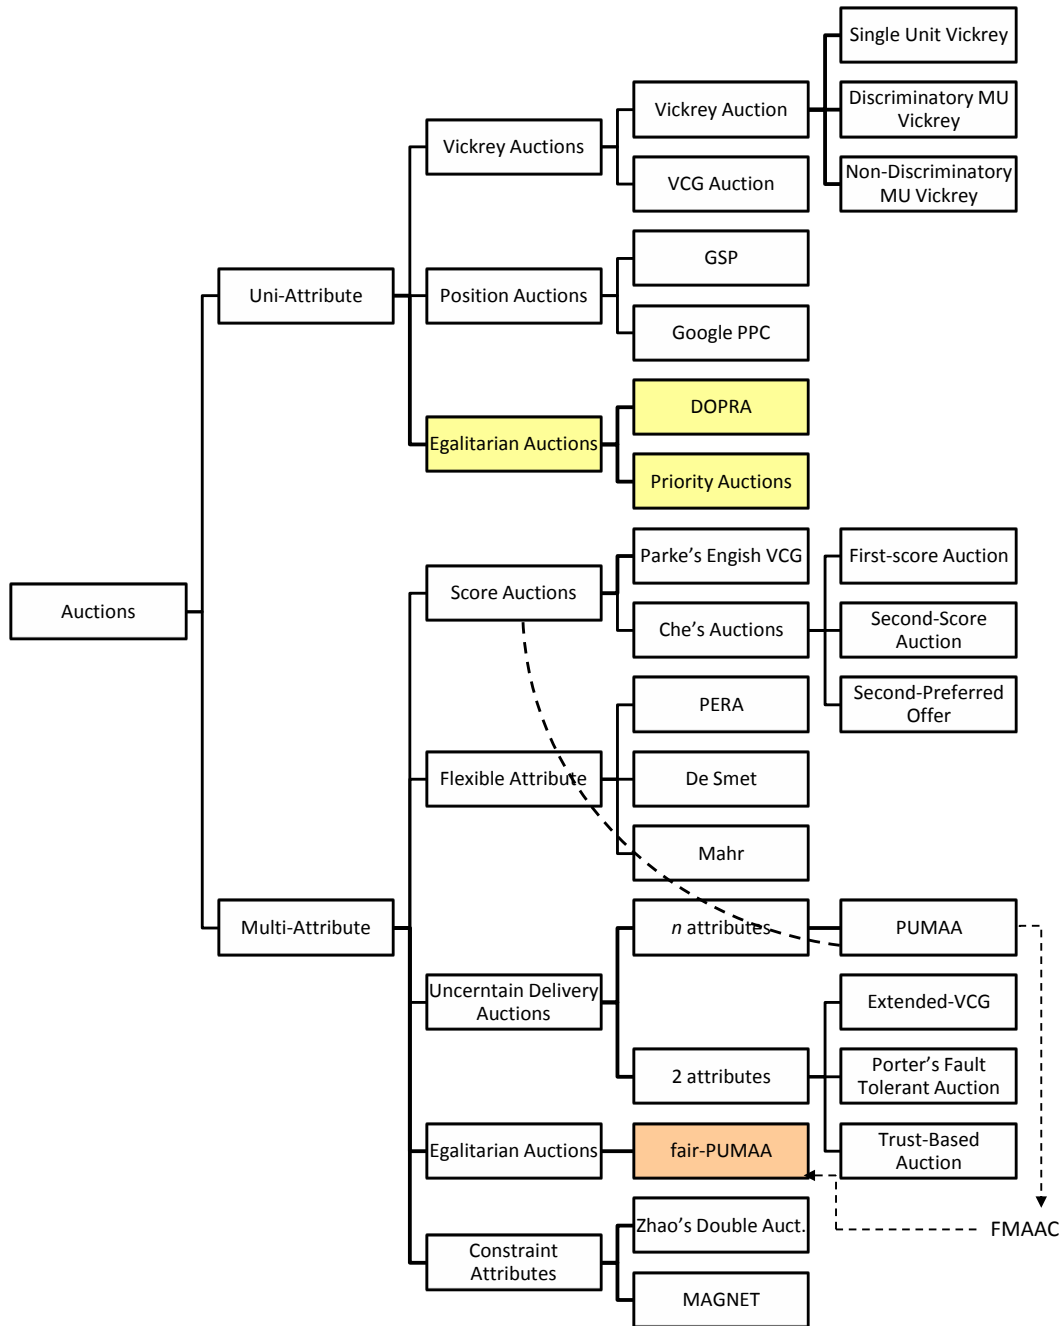


Figure 5.1: fair-PUMAA and other egalitarian auctions classification with respect to the auctions described in 2.

The steps followed in fair-PUMAA are as follows:

1. **Call for proposals:** The auctioneer a_0 defines the task T_0 to be auctioned and the set of attributes which conditions the utility (e.g. delivery time, minimum quality, etc.). In this way the auction is defined within the following settings:
 - Unverifiable bidder-provided attribute: The economic amount b_i which a bidder a_i asks for performing the task.
 - Verifiable bidder-provided attributes: The physical attributes related to the task which the auctioneer evaluates (e.g. delivery time) $A_i^v = \langle at_{1i}^v, \dots, at_{mi}^v \rangle$.
 - Auctioneer-provided attribute: The priority $w_i \in [0, 1]$ corresponding to the bidder a_i where 0 means that there is no risk of a_1 leaving the market and 1 means that there is a high risk.
2. **Bidding:** Interested bidders provide bids B_i , which are composed of the price b_i and the values of the requested attributes, $at_{1i}^v, \dots, at_{mi}^v$, such that $B_i = \langle b_i, at_{1i}^v, \dots, at_{mi}^v \rangle$. Given that the bidder must evaluate if they really want to participate into the auction, as their participation can condition their priority.
3. **Winning determination problem:** The auctioneer receives the bids, and extends them with the bidder priorities, $B_i' = B_i \oplus \langle w_i \rangle$. It then selects the best bid according to an aggregation function V_0 :

$$\operatorname{argmin}_i (V(B_i')) \quad (5.1)$$

As an example we propose:

$$V_0(b_i, A_i^v, w_i) = \frac{b_i * \prod_1^m at_{mi}^v}{1 + w_i} \quad (5.2)$$

4. **Payment mechanism:** The payment follows PUMAA, treating the priority as another attribute:

$$p_1 = \begin{cases} V_0^{-1}(V_0(b_2, A_2^v, w_2), A_1^v, w_1) & \text{if } A_1^v \succeq A_2^v \\ V_0^{-1}(V_0(b_1, A_1^v, w_1), A_1^v, w_1) & \text{if } A_1^v \prec A_2^v \end{cases} \quad (5.3)$$

Which given the example of Equation 5.2 is:




Figure 5.2: Priority classification methods analyzed according to the measure used (auctions won or auctions lost), the bid evaluation type (quantitative or qualitative) and the update process (deterministic or stochastic).

$$p_1 = \begin{cases} \frac{b_2 * \prod_1^m at_m^v * (1+w_1)}{1+w_2 * \prod_0^m at_m^v} & \text{if } A_1^v \succeq A_1^v \\ \frac{b_1 * \prod_1^m at_m^v * (1+w_1)}{1+w_1 * \prod_0^m at_m^v} & \text{if } A_1^v \prec A_1^v \end{cases} \quad (5.4)$$

5. **Attribute update:** The auctioneer updates the information it has concerning the bidders. In this case, it computes the priority corresponding to each bidder.

5.4 Priority Calculation Methods

When defining the methods for computing priority, one option is to directly relate the priority to the number of victories and defeats bidders have obtained in auctions in which they have participated. With this approach, the dimensionality of the auction is directly represented in the priority, as the aggregation of attributes is implicit in the winner determination problem (defined by an aggregation function involving all the auction attributes weighted by the auctioneer's preferences). This method offers a quantitative approach which takes into account if a bidder has won or lost. An alternative is to combine information regarding victories with the bids offered by bidders and the evaluation function used by the auctioneer. This approach

		Timestamps	0	1	2	3	4	5	6	7	8
a_1	Bid Sequence	-	×	×	✓	×	×	✓	×	×	
	WOC w_1	0.00	0.50	0.66	0.50	0.60	0.66	0.57	0.63	0.66	
	LS(8) w_1	0.00	0.13	0.25	0.00	0.13	0.25	0.00	0.13	0.25	
a_2	Bid Sequence	-	✓	✓	×	×	×	×	×	×	
	WOC w_2	0.00	0.00	0.00	0.25	0.40	0.50	0.57	0.63	0.66	
	LS(8) w_2	0.00	0.00	0.00	0.13	0.25	0.38	0.50	0.63	0.75	
a_3	Bid Sequence	-	×	×	×	✓	✓	×	✓	×	
	WOC w_3	0.00	0.50	0.66	0.75	0.60	0.50	0.57	0.50	0.55	
	LS(8) w_3	0.00	0.13	0.25	0.375	0.00	0.00	0.13	0.00	0.13	
a_4	Bid Sequence	-	×	×	×	×	×	×	×	✓	
	WOC w_4	0.00	0.50	0.66	0.75	0.80	0.83	0.86	0.88	0.77	
	LS(8) w_4	0.00	0.13	0.25	0.375	0.50	0.63	0.75	0.88	0.00	

Table 5.1: An illustrative example of the auction record of three different agents and the priority they would obtain when using the won auction and the losing streak coefficient (with a maximum losing streak threshold ml of 8). ✓ means that the bidder has won the auction whilst × means it has not.

allows to evaluate the priority in a qualitative way as the bid information can be used to determine how far from victory a bid was. In this way higher priority can be given to the bidders who do not win but offered high adjusted bids. Finally, both the quantitative and qualitative methods can be also modeled under an stochastic model, this offsets some possible side effects of fairness mechanisms (see Figure 5.2).

5.4.1 Quantitative priority methods

Quantitative priority methods take into account the number of victories and defeats achieved by each agent. Using this information the auctioneer gives more priority to bidders with a high defeat record, reducing the risk of these agents leaving the market. Two base methods have been defined: won auction coefficient (WOC) and losing streak (LS).

Won auction coefficient (WOC).

The won auction coefficient is the most intuitive way of computing the priority of an agent. It establishes the ratio between the number of auctions in which a bidder has participated

($par(a_i)$) and the number of auctions in which it has won ($won(a_i)$). As shown in the following expression:

$$w_i = 1 - \frac{1 + won(a_i)}{1 + par(a_i)} \quad (5.5)$$

This measure gives the auctioneer an idea of the proportion of won auctions for each bidder since the auctioneer's market entry. Using this priority, a bidder which has won all the auctions it entered will have a priority of $w = 0$ whilst a bidder which has never won an auction will have a priority close to 1.

To illustrate the behavior of WOC consider, for example, the sequence of wins (\checkmark) and defeats (\times) obtained by the set of agents shown in Table 5.1. The WOC method takes into account information regarding the entire period where an auctioneer has been active. Thus, in the long run the auctioneer will tend to compensate bidders who have lost many times during initial auctions and it will ballast those which have won the most. This means that if a bidder wins an auction after a large period of defeats, its priority will remain high as it still has a relatively low ratio of victories. In Table 5.1 example, a_4 won the most recent auction but it still has the highest WOC priority because it is the agent which won the least number auctions. A bidder experiencing a long losing streak will have a low priority if he obtained a higher number of victories in the past. For instance, in Table 5.1, a_2 has had a long losing streak but, at the end, it has the same priority as a_1 due to the fact that they have both won 2 auctions (a_2 at the very beginning, a_1 in a more evenly spaced way). A time window can be adopted to prevent such a scenario (e.g. computing only the auctions summoned during the last week); similarly, limiting the auction record which is taken into account (e.g. computing only the last 100 auctions summoned by the auctioneer) could alleviate the side effects of WOC caused by its long term memory.

Losing streak (LS).

Under certain circumstances, especially when some of the bidders are directly controlled by human entities, certain bidders may experience cognitive distortions which can affect their reasoning [51]. A human-controlled bidder may be more susceptible to leaving the market after a sequence of two wins and 6 defeats (as agent a_2 of Table 5.1 with a WOC of 0.66) rather than after a succession with the same amount of won auctions but where the victories are more spread (e.g. the case of agent a_1 with 2 defeats, 1 victory, 2d, 1v, 2d).

In these scenarios priority can be defined as the basis of the bidders losing streak $ls(a_i)$: the number of consecutive auctions which an agent a_i has lost. The longer the losing streak

of an agent, the higher its priority is. For that purpose, we define a tolerance threshold ml (maximum losing streak) such that after losing ml consecutive auctions, a bidder is as susceptible to leave the market as one who has lost $ml + n$ auctions. Each auctioneer fixes ml according to its beliefs about bidders behaviors. As the auctioneer does not know the different tolerances (to losing streaks) of the different bidders, it must estimate a generic ml which will be the same for every bidder. Moreover, using a different ml for each bidder leads to an unjust starting point as not all bidders would be treated under the same rules. For example if two agents exchange their bid sequences they may obtain different results as they are participating under different conditions. Therefore, the same ml is assigned to each bidder.

We use this threshold to define the highest priority of the bidder as follows:

$$w_i = 1 - \frac{\max(0, ml - ls(a_i))}{ml} \quad (5.6)$$

where $ls_i(a_i)$ is the losing streak of bidder a_i . In this way, a bidder that has just won an auction has a priority of $w = 0$ whilst a bidder which has gone ml auctions without winning has a priority of $w = 1$. Unlike the previous method, we can say that the LS method has no long term memory as a bidder will see its priority reduced to 0 after winning an auction, independently of its results in the past.

Following the example in Table 5.1, we can see how despite agent a_1 and a_2 having the same number of victories and same priority when using WOC, a_2 has a higher priority when it is computed using LS due to the fact that it has the highest losing streak.

5.4.2 Qualitative bid-based priorities

The methods presented above have an absolute behavior, since they treat all non-winning bidders in the same way without taking into account the attribute values provided in their bids. In other words, the priority method treats the bidder who almost won the auction (the second best bid) and the bidder who offered the worst bid as the same. On the one hand, this situation may encourage some bidders to provide dummy bids with no chance of winning in order to increase their priority, obtaining a higher chance of winning one of the following auctions. On the other hand, some bidders offering almost-winning bids could feel frustrated by the fact that attribute values of bids have no effect on the priority outcome. In order to avoid this situation, we propose taking into account the attribute values of bids offered in the priority definition.

For this purpose, we define the fitness q_i of a bid B'_i (which includes all the attributes involved in the auction) in relation to the winner bid of the auction B'_1 depending on whether the auction

	Timestamps	0	1	2	3	4	5	6	7	8
a_1	Bids evaluation	-	7	10	8	6	5	10	9	9
	Bid fitness q_1	-	0.875	1.000	0.888	0.856	0.555	1.000	1.000	1.000
	bbWOC w_1	0.000	0.467	0.304	0.469	0.567	0.614	0.514	0.443	0.388
	bbLS w_1	0.000	0.109	0.000	0.111	0.218	0.288	0.000	0.000	0.000
a_2	Bids evaluation	-	7	8	9	5	9	9	8	7
	Bid fitness q_2 A2	-	0.875	0.800	1.000	0.714	1.000	0.900	0.889	0.778
	bbWOC w_2	0.000	0.467	0.626	0.456	0.544	0.443	0.523	0.582	0.623
	bbLS w_2	0.000	0.109	0.209	0.000	0.089	0.000	0.113	0.224	0.321
a_3	Bids evaluation	-	8	5	7	7	7	6	5	4
	Bid fitness q_3	-	1.000	0.500	0.778	1.000	0.778	0.600	0.556	0.444
	bbWOC w_3	0.000	0.000	0.200	0.390	0.299	0.407	0.470	0.517	0.549
	bbLS w_3	0.000	0.000	0.063	0.160	0.000	0.097	0.172	0.242	0.297

Table 5.2: An illustrative example of the auction record of three different agents, the evaluation received by their bids. Examples of the priority they would obtain when using the bid based won auction coefficient and the bid based losing streak (with a ml of 8). Winning bids are marked in bold.

is reverse (left) or not (right):

$$q_i = \frac{V_0(B'_i)}{V_0(B_i)} \quad q_i = \frac{V_0(B'_i)}{V_0(B'_i)} \quad (5.7)$$

Following this definition we can say that the winner bid of the auction has the maximum fitness $q_1 = 1$ and that the higher q_i is for a bid, the closer it will have been to win the auction. Thus, this fitness metric provides a proportional measure of how far from the auction winner a bid was, no matter the magnitude of the bids. Given that the fitness of a bid is always a value between 0 and 1, this measure can be used to compare results from different auctions even if the bids are of different magnitude or if they involve different attributes.

Taking bid fitness into account, we propose two new methods, based upon the *Won Auction Coefficient* and the *Losing Streak* priority calculation methods presented above.

Bid-based won auction coefficient (BBWOC).

In this mechanism the priority is calculated in the same way as in WOC method. Instead of taking into account the number of auctions in which an agent has participated, it gathers the fitness of the bids it offered in the j auctions where it has participated.

$$w_i = 1 - \frac{1 + \text{won}(a_i)}{1 + \sum_{j=0}^{c-1} q_i^j} \quad (5.8)$$

where q_i^j is the fitness of the bid B_i^j in the j auction and c indicates the current auction.

In this way, if there are two bidders who have won the same number of auctions, the bidder which has provided higher quality bids without winning will be the one with the highest priority. This can be observed in the example provided in Table 5.2 where bidders a_2 and a_3 have won the same number of auctions, however, at the end, a_2 has a higher priority (0.623) than a_3 (0.549) as a_2 have been providing better bids.

Bid-based losing streak (BBLS).

The same action can be performed with the LS method. This method follows the priority function in Equation 5.6. Instead of using the losing streak accumulated by bidders to define the priority, we propose using the summation of the bid fitness obtained during the losing streak. In this way a long streak of almost winning bids will increase the priority faster than a long losing streak of low quality bids.

$$w_i = 1 - \frac{\max(0, ml - \sum_{j=c-ls(a_i)}^{c-1} q_i^j)}{ml} \quad (5.9)$$

In Table 5.2 it can be seen that, despite a_2 having a shorter losing streak than a_3 (3 consecutive loses against 4), a_2 ends with a higher priority (0.321 against 0.297) due the fact that its bids have been closer to the victory than the ones bid by a_3 .

5.4.3 Stochastic priorities.

A problem that may arise due to priorities is the alteration of the agents wealth ranking. For example the richest agent when not using the fairness mechanism could become the second richest agent when priorities are used (see experiment described in Section 6.4). This problem can be minimized by updating the priority of the bidder following stochastic or probabilistic methods, so sometimes the wealth rank is altered and some times it is not. In addition, a probabilistic approach may prevent agents from learning the priority mechanism. Thus, step 5 (attribute information update) of the protocol is modified so the bidders priority and the information regarding the auctioneer provided attributes (number of won auctions, participated auctions etc.) is updated according to an update priority probability parameter $up_i \in [0, 1]$

defined by the auctioneer a_i . The priority of the bidder's can be computed using any of the previously presented methods, the difference is that with this stochastic approach the priority is not always updated.

In Algorithm 5.1 we present a stochastic algorithm which can use equations 5.5, 5.6, 5.8 and 5.9 as functions to calculate the priority of bidders. Depending on the equation chosen to determine the probability we can refer to the method as probWOC, probLS, probBBWOC or probBLS.

Algorithm 5.1 Stochastic priority algorithm (auction step 5)

```

1: if  $Random(0, 1) < up_0$  then
2:   for all Bidder  $a_i \in$  auction_participants do
3:     update  $a_0$  auction_statistics
4:     update  $a_0$ .priority[ $a_i$ ]
5:   end for
6: end if

```

Using a stochastic method bidders will not have the certainty of which their priorities, thus, agents will have more difficulties in taking advantage of learning their priorities. Moreover, the use of an update probability can smooth the alteration of the wealth ranking as it is reasonable to expect that a low update priority probability will reduce the fairness of the final allocation. It will also reduce the wealth ranking. Therefore this approach raises the issue of finding a good compromise between fairness and the preservation of the wealth ranking.

5.5 Considerations Regarding fair-PUMAA Properties

The previous chapter states that the use of auctioneer provided attributes in FMAAC may affect the auction properties. Therefore, in this section we analyze how the use of priorities affects fair-PUMAA properties.

There are five properties which are not altered due to the use of priorities: fair-PUMAA does not punish or charge a fee for participating in the auctions, thus the mechanism remains **individual-rational**; the transfer between players at the end of an auction is still equal to 0, respecting **budget balance**; fair-PUMAA still preserves the auctioneer expected utility in the case of a faulty delivery task and punishes agents which do not respect their offered attributes. Therefore the mechanism remains in the same degree of **robustness and reliability** as PUMAA. Finally, given its second-price payment philosophy, fair-PUMAA is **not buyer-optimal**.

On the other hand, there are three properties which are affected by the use of priorities:

- The first and most obvious condition affected by the use of priorities is the **social welfare** of the resulting allocations. PUMAA produces utilitarian allocations, maximizing the agents' utility without considering any other aspect. Conversely, fair-PUMAA produces more **egalitarian** allocations as the use of priorities gives agents more chances of winning an auction. In comparison to those agents with bad auction records, reducing the difference between richest and poorest agents.
- According to the **efficiency** definition in [8], fair-PUMAA is no longer efficient in a single-shot auction. A bidder with a bid B_i with a price and a bundle of attributes which are not the best ones (a dominated bid) could win the auction due to the inclusion of priorities. However, if we consider the bid as the extended bid B'_i (which includes the priority) the winner bid will not be dominated. Thus, efficiency in fair-PUMAA is conditioned to the definition of what is the bid: the original bid B_i provided by the bidder or the auctioneer extended bid B'_i .

Moreover, if we consider the efficiency definition of [42], experimentally we can see how, under certain circumstances (e.g. when bidder coalitions try to remove agents from the market), fair-PUMAA can be as efficient (or even more efficient) than PUMAA as at the end of a long auction sequence players can obtain similar or higher utilities (see the experiments of the next Chapter).

- Finally, **incentive compatibility** may be affected due by the use of priorities. In a single shot-auction the mechanism remains incentive compatible (section 3.4 is still valid in this scenario). On the other hand, if we consider that the result of an auction will affect the chances of a bidder of winning the next one, there could be cases where a bidder decides to lose on purpose, in order to increase their priority. This situation can be minimized depending on the priority method used, but cannot be completely avoided. Thus, in fair-PUMAA, incentive compatibility is sacrificed for the sake of equity.

This bleach can be used by bidders to increase their chances of winning a hypothetical future auction however, conversely to priority auctions [54], this mechanism cannot be used to reduce the market price as in a single-shot auction the maximum utility for a bidder will be obtained by bidding truthfully. Therefore, in fair-PUMAA, if a bidder lies regarding any of their attributes (offering worse than true) it will be to intentionally lose an auction the bidder was not interested in. Thus, the auctioneer does not lose negotiation power as bidders cannot decrease the price of the market by lying in a single-

shot auction.

5.6 Summary

The bidder drop problem is a common issue which auction designers face in repetitive and recurrent auctions. Fairness mechanisms have proven useful in reducing this problem in uni-attribute auctions. In this chapter we proposed a multi-dimensional fairness mechanism for the multi-attribute auction domain.

Fairness mechanisms for multi-attribute resource allocation problems, analyzed from a multi-dimensional perspective, need to take into account all the attributes which affect the allocation. This is due to the fact that focusing fairness only on the revenues of the participants may compromise the quality of the rest of the attributes involved in the decision process (e.g. delays in deliveries, alterations in energy consumptions, etc.). In this chapter we propose a multi-dimensional fairness approach for multi-attribute auction resource allocation based upon the definition of priorities. Unlike existing priority auctions which follow a first-price policy [54], our proposal faces the problem in a more holistic way. The FMAAC structure allows one to take into account priorities during the entire auction process (in WDP but also in the payment mechanism) and to define a second-price based payment method which reduces the asymmetric balance of negotiation power.

The proposed approach assigns higher priorities to the bidders most likely to leave the auction market (due to their bad results in previous auctions). With this method the agents most likely to leave the market have most chance of winning an auction. This method gives these agents an incentive to remain in the market. Thus, the process of computing priorities becomes a key factor in fair-PUMAA.

We presented a collection of eight methods for computing the priority of the agents based on the ratio of auctions won or lost, or their losing streaks. Using the results of the auction to compute the agent's priorities implicitly means to take into account all the auction attributes since the auction winner is determined by the dimensionality of the WDP. Two of the methods presented were designed under a quantitative schema (absolute values regarding auctions won or lost). Two others were designed under a qualitative schema, giving higher priorities to those bidders which offered good bids but lost the auction rather than to bidders which offered low quality bids. A stochastic version of these four methods is also provided. The stochastic approaches proposed allow the tuning of the consequences of the priority methods regarding the wealth ranking disorder.

EXPERIMENTATION AND RESULTS

This chapter presents the tests and results of the methods developed during the course of this thesis. The experiments in this chapter have been performed in a multi-agent based auction simulator in which agents compete to sell and to buy tasks, following a typical schema for supply chain management.

In the first section we briefly introduce the goals of the experiments performed. Then, before presenting the results of the research conducted we introduce the simulation environment and the data used. In Section 6.3 we present the results of the Preserving Utility Multi-Attribute Auctions (PUMAA) described in Chapter 3. Finally, in Section 6.4, we test the multi-dimensional fairness mechanism fair-PUMAA (Chapter 5) which has been developed via the Framework for Multi-Attribute Auction Customization (FMAAC) of Chapter 4.

6.1 Introduction

This chapter presents seven experiments, four corresponding to the evaluation of PUMAA and three corresponding to the evaluation of FMAAC and fair-PUMAA (see Table 6.1).

The first PUMAA experiment the need to use multi-attribute auctions in auction scenarios where more than one attribute can influence the quality of allocations. The goal of the second experiment is to empirically support the theoretical incentive compatibility demonstrations performed in Chapter 3. The third experiment concerns the definition of PUMAA's evaluation function and how the evaluation function conditions the allocations. The fourth experiment studies the robustness of PUMAA against cheating agents and analyzes how an auctioneer can preserve its utility when tasks are not delivered satisfactorily.

The other three experiments assessed in this chapter investigate the effectiveness of FMAAC

	Experiment	Goal
PUMAA	1	To study the need for multi-attribute auctions
	2	To analyze the strategy-proofness of PUMAA
	3	To evaluate the impact of the evaluation-function
	4	To evaluate PUMAA's robustness against cheating To study PUMAA's utility preservation
fair PUMAA	1	To compare uni-dimensional and multi-dimensional fairness To compare the different priority methods
	2	To analyze the behavior of stochastic priorities
	3	To study the way multi-dimensional fairness affects BDP

Table 6.1: Summary of the experiments performed.

in the implementation of the fair-PUMAA mechanism. The goal of the first experiment is to compare the allocations obtained when using uni-dimensional or multi-dimensional priorities, moreover this experiment compares the outcomes obtained by different priority calculation methods. The second experiment analyzes the influence of the update probability parameter in the stochastic priority methods. Finally, the last experiment studies the way in which the multi-dimensional fairness can be used to reduce the bidder drop problem.

6.2 Simulation Environment and Data

To test and illustrate the behavior of the mechanisms presented, we simulated the service allocation process of an industrial environment using a multi-agent system. In this section we describe the simulation tool developed for this purpose and the kind of data used for the simulations.

6.2.1 Multi-agent simulation

The simulation tool used to test the mechanisms described reproduces the operation of an industrial environment where companies need to solve different issues [63] (e.g. unforeseen faults, local system crashes, customer problems, etc.) by outsourcing certain tasks. When an incident is detected, a preliminary diagnosis process determines the incident typology, the task (or set of tasks) which must be carried in order to solve the incidence and its maximum deadline. It, then, determines the order in which the tasks must be performed and the tasks




Figure 6.1: Multi-agent system architecture utilized within the simulation

which need to be assigned to an external provider.

In the simulation, each entity (which can correspond to a company, a team, a department, etc.) is represented by an agent which acts selfishly, working solely for its own benefit. The participants in the simulation can play two different roles: they can act as service agent *SA* (an agent which controls a workflow and needs to externalize a task) or as a resource provider agent *RP* which wants to perform a task for a payment. Despite the fact an entity can perform both roles (but not within the same auction), for the sake of simplicity and to make the experiments more comprehensible, agents can take only one role during the simulation. Thus, service agents will always act as the auctioneers whilst resource provider agents will always be bidders (Figure 6.1).

- **Service Agents:** When a service agent detects an incidence, a preliminary diagnosis process determines its typology, its maximum delivery time and other possible attributes required. Then, the agent calls for an auction in order to outsource the task which will resolve the incidence. If none of the resource providers are available to carry out the task at the time of the auction, the SA repeats the auction after a certain period of time (between 5 and 10 time instants). In case of not finding an agent to develop the task, the service agent can also opt for relaxing the conditions of the auction (e.g. extending the deadline of the task) assuming, in this way, a quality reduction (e.g. a delay).

Each task is defined, at least, by its typology (which determines the type of resources required to perform the task) and its delivery time.

- **Resource Provider Agents:** Each resource agent is defined by its capacity (the number of concurrent tasks it can perform), its typology (which determines which types of tasks it can perform), the time distributions that define how long it takes to finish each type of

task and its true attribute values (the cost it incurs when performing a task with a given set of attributes).

Thus, when a resource provider receives a call for proposals it evaluates whether it can perform the task required (given its capacity and type) and its true values (its cost and expected delivery time). With this information the agent can decide if it wants to participate into the auction and defines the bid to be submitted. The bid it submits can be truthful (with true values) or untruthful. However, the simulation does not allow bidders to lie regarding their identity or typology (in non-simulated environments these aspects can be controlled by credential systems [18]).

The simulation follows an event-driven structure. Every time something happens (e.g. an auction is started, an item is delivered, a payment is performed, etc.) an event is added to a time-ordered event queue. Each event contains information concerning the action to be triggered, the time instant in which the action must be performed, the agents which are involved in the event and other attributes which depend on the event type. Thus, agents will keep inserting events to the event queue, whilst the system executes the actions associated with each event in a sequential way.

6.2.2 Data used in the simulations

Two different types of data have been used in the simulation. The first one corresponds to a simple synthetic example which can be used to analyse the behavior of the mechanisms whilst the second one is based in real data extracted from an industrial environment.

Synthetic data

We define three sequential workflows (managed by three different agents $SA1$, $SA2$ and $SA3$) which are composed of six different tasks and which have a maximum workflow duration of 90 time units (Figure 6.2). Each task of the workflows needs a given type of resource (between A and D) in order to be performed. Thus, each workflow requires between 1 and 4 different resource types to be performed. The start of a workflow is defined by a probability of p_{SA} for each time unit, meaning that at a given time instant there is a probability of p_{SAX} of starting a workflow managed by the Service Agent X . This probability can be varied for each experiment. Moreover, the type of resource which each task needs to be outsourced is manually or randomly defined before an execution (in order to obtain significant data, an execution can have n simulations of the same experiment).




Figure 6.2: Example of the types of resource required by each service agent.

	RP1			RP2			...	RPn		
	T	t	c	T	t	c	...	T	t	c
Skills	A	15	25	B	15	24	...	D	13	26
	C	14	26	D	13	25	...	B	13	26
	B	10	30	A	10	31	...	A	12	28

Table 6.2: Example of resource skills (T = type, t task duration) and costs (c) randomly generated per each agent.

The number of resource providers participating in each experiment can be set manually. Each resource provider agent has the skill to perform three different tasks (each resource has three types) with different qualifications of cost and delivery time (see example in Table 6.2). The type of tasks that a resource provider can perform (between A and D) and the skills for performing each type of task (a duration between 10 and 15 time units and a cost between 20 and 30€) are set before each experiment execution (manually or randomly, depending on the experiment goals).

Production factory data

The second type of data used for the experiment is based upon real data extracted from an industrial organization ¹. The data contains information regarding unexpected issues (unexpected faults, customer complaints, revisions, etc.) the company has had to deal with. Each

¹Due to confidentiality agreements we cannot reveal the name of the company

Type of task	S1	S2	S3	S4	S5	S6	S7	Total
Interventions	1,604	3,892	868	3,256	1,400	2,120	1,052	1,4192
Proportion	11.30%	27.42%	6.11%	22.9%	9.86%	14.93%	07.41%	100%

Table 6.3: Number of instances of each type of task in the original data.




Figure 6.3: Proportion of each type of task in the original data.

unexpected issue entailed the execution of a single independent task which covered the whole process of solving the problem. The data was collected over two years and is composed of information concerning the tasks which were performed in order to solve these issues (the type of task performed, the emergency level of each task, tasks deadlines, the partner or team which developed each task, task durations, etc.). During this period occurred 14,192 interventions of seven different typologies (see Table 6.3 and Figure 6.3).

Each type of task can only be developed by certain resource providers (see Table 6.4). The duration of each task is conditioned by the type of task and for the resource provider which is performing it. Regarding the deadline of each task, it is determined by the emergency level parameter associated to each particular task (see Table 6.5). The emergency level is independent from the task typology but it is related to the type of customer who suffers the incidence (for instance, a fault in a hospital has a higher emergency level than a fault in a private home). Using this information we have built a dataset where service agents are in charge of managing 7 different types of independent tasks and where 8 types of resource providers compete to obtain those tasks. Using this data we have estimated the probability distribution function that




Figure 6.4: Left) Probability of a task of type S3 occurring in a given time instant. Right) Probability function that defines the emergency level of a task of type S3.

	S1	S2	S3	S4	S5	S6	S7		S1	S2	S3	S4	S5	S6	S7
RP1	✓	✓	✓	✓	✓	✓	✓	RP5			✓		✓		✓
RP2	✓	✓	✓	✓	✓	✓	✓	RP6					✓		
RP3	✓	✓	✓	✓	✓	✓	✓	RP7		✓	✓		✓		✓
RP4	✓		✓	✓	✓	✓	✓	RP8		✓	✓		✓		

Table 6.4: Relation of the tasks S that a resource of a type R can perform.

represents the task occurrence during an average day (see Figure 6.4). The same has been done to model the urgency parameter distribution and the task service durations of each resource provider (each resource provider has different time distributions for each kind of task, see Figure 6.5). These distributions are then used to execute the simulations².

Thus, in experiments which use the production factory data set the agents are defined in the following way:

Each service agent is in charge of only one type of task and is defined by the following parameters:

- The type of the service task of which it is in charge (1 to 7). Each type of task can only be performed by certain types of resource provider (see Table 6.4).
- The probability distribution function that defines the task occurrence. In the simulation, tasks appear following this probability function.

²The specific probability distributions for each type of resource and task can be found in [81].




Figure 6.5: Distribution functions that determine the duration of each type of task for resource provider RP7

Emergency Level	1	2	3	4
Maximum duration (minutes)	30	120	300	600

Table 6.5: Relationship between the urgency of a service and its maximum execution time.

- The probability distribution function that defines the emergency level of the extant tasks. Every time the service agent wishes to outsource a task it assigns an emergency level to it according to its emergency parameter distribution function. The emergency level is used to establish the execution deadline for each task (see Table 6.5).

On the other hand, each resource provider agent defends the interests of an outsourcing company. Each type of resource provider agent is defined by the following parameters:

- The type of the resource provider (1 to 9). The type of the resource provider determines the tasks which the agent can perform (see Table 6.4).
- Its capacity. If nothing is specified, we assume that each resource provider can perform only one task at a time.
- A probability distribution function specifying the time it takes to perform each task. If the resource provider can perform n tasks, the resource provider will have n different

time distributions.

- The cost the agent incurs for performing each task.
- Its bidding strategy (which can be adaptive bidding, honest bidding or dishonest).

Bidders can follow three bidding strategies. They can bid honestly, meaning that they offer their real economic cost and the delivery time they consider they will provide. Bidders can also follow an adaptive strategy, adapting their offers according to the resulting allocations in order to increase the chances of winning the auction and maximizing their benefits [43]. In this case, bidders start by offering their true values and, when they win some auctions, they slightly increase their economic ambitions to try to increase their revenue. If they cease to win auctions after rising its cost attribute, bidders gradually reduce their ambitions until they bid their actual economic cost. Finally, agents can follow a cheating strategy. In such situation bidders always offer values different from their true-values (either higher or lower) both in term of cost and delivery time. Dishonest bidders keep track of the incomes they have obtained and the bid configuration they have used every time they have won an auction (e.g. selling a product 10% more expensive and offering a smaller delivery time than the actual reported them a revenue of 10). Using this information, they decide whether they will try a new bid configuration, or if they will follow a bid configuration that has brought them a good utility in past [64].

In most of the experiments where the production factory dataset has been used, there is one service agent for each task (7 SAs) and one for each type of resource provider (8 RPs). If this is not specified in the experiment description, this configuration is set as the default configuration for the experiments which use this data.

6.3 PUMAA Results

This section presents the experimentation and tests performed in the PUMAA mechanism described in Chapter 3. We have performed 4 different experiments with four differentiated goals. The first experiment compares the resource allocations obtained when only one attribute is taken into account in the allocation decision process, with the allocations obtained in multi-attribute auctions. The second experiment evaluates the impact which the evaluation function has upon the allocations. The third experiment studies the incentive compatibility of PUMAA. The final experiment compares PUMAA's performance with other multi-attribute auction mechanisms in terms of robustness.

6.3.1 Experiment 1: The need of a multi-attribute auction

The goal of this experiment is to compare the results obtained when only one attribute is taken into account during the allocation process, or when all the attributes involved in the problem are evaluated together. For this purpose we compare the allocations obtained when service agents allocate tasks using uni-attribute auctions (first-price and Vickrey), when they choose the fastest resource and when using a multi-attribute auction (PUMAA). The experiment is performed using the production factory dataset. In this experiment we do not expect that any mechanism will overcome any other, our goal is to analyze the tradeoff among them.

Scenario

This scenario analyzes the performance of PUMAA, a multi-attribute method, against three uni-attribute allocation methods terms of economic cost and delays. In this scenario 7 service agents (one for each type of task within the data set) auction tasks whilst eight resource provider agents (one for each type of resource in the data set) compete to perform the tasks.

The experiment is performed four times using different types of allocation methods. The simulation time corresponds to 180 time cycles (or days). Each simulation is repeated 200 times in order to obtain significant data:

The allocation methods used in this experiment are the following:

- **First price uni-attribute auction:** The resource allocation is performed using a first-price auction where the resource providers bid their economic cost. The auction winner is determined by the lowest bid. Thus, the available resource provider with the cheapest price is in charge of performing the service. The payment for the winning bidders corresponds to the economic amount indicated in the submitted bid.
- **Vickrey uni-attribute auction:** The allocations is solved using a Vickrey auction where available bidders submit their economic costs. The auction winner is the one with the best bid and it's payment corresponds to the second best bid.
- **Best delivery time:** The resource provider with the fastest delivery time is the agent which will carry out the service. Its payment value corresponds to the economic amount indicated in the winner's offer.
- **PUMAA:** The allocation is determined using the PUMAA mechanism, with the product as evaluation function ($V_0(b_i, t_i) = b_i * t_i$).




Figure 6.6: Total task allocation cost for each Service Agent using different allocation mechanisms: First-price auction (green), Vickrey auction (clear blue), choosing the fastest resource provider (purple) and PUMAA (red)

In this experiment all the bidders follow an adaptive bidding strategy: in order to increase their chances of winning an auction and to maximize their utility bidders can modify the economic amount they submit depending if they have won or lost the past auctions [43].

The results of this experiment are analyzed in terms of the total expenditure that service agents had to spend to externalize the services (the lower the better). The mean service externalization cost and the number of delays (tasks exceeding their deadline) produced on the task executions using the different types of allocation methods (the lower the better) are also taken into account.

Results

Figure 6.6 shows the mean economic amount, per simulation, which the 7 different service agents had to spend in order to allocate all their tasks to a resource provider. Green (first-price auction) and clear blue (Vickrey auction) bars show the resulting economic costs when following an approach based only upon the economic attribute. Purple bars show the costs when the task is assigned to the fastest resource, whilst the red bars corresponds to the results obtained when PUMAA has been used.




Figure 6.7: Mean cost for a service allocation when using different allocation methods.

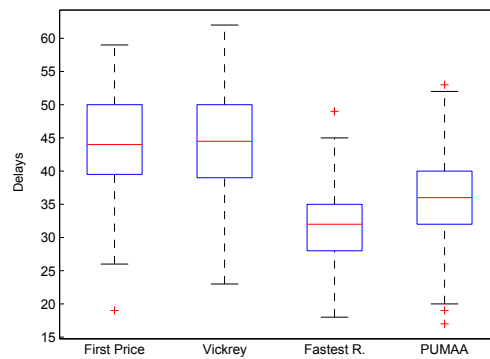


Figure 6.8: Delays produced in the simulation when using different allocation methods.

The bar chart shows that following the strategy of picking the fastest resource provider leads to a higher economic cost for service execution. This is a logical result as this strategy ignores the economic cost of the resources. When observing the other three allocation strategies it can be seen that for all service agents but SA3 the most expensive allocation strategy is to use the first-price attribute. This situation is produced by the lack of incentive compatibility for first-price auctions, where some agents may have obtained higher utility by asking higher prices than their true-values. If we compare PUMAA with the Vickrey auction we can see that they obtain similar costs: for agents SA3, SA4 and SA6 the Vickrey auction is slightly cheaper whilst for the rest PUMAA represents the most economic choice.

These observations are confirmed in Figure 6.7 which shows a box plot of the mean cost of a service in the different simulations; offering a global view of a service cost from the service agents perspective. This figure points that the fastest resource picking and the first price auction provoke higher task costs than the Vickrey auction and PUMAA. The similar costs generated by PUMAA and the Vickrey auction can also be observed in this boxplot as there are no significant differences between their two boxes.

Regarding the delays produced in the system Figure 6.8 shows a boxplot with the number of delays which have been produced with each allocation method. In the boxplot we can see how the allocation strategy which obtained a lower number of delays is the fastest resource choice. The worst outcome was achieved by the strategies based upon price. This is not surprising as the first policy is focused exclusively on time efficiency whilst the other uni-dimensional strategies ignore this parameter. If we focus on PUMAA we can see that, despite its results outperforming the ones obtained with the fastest resource strategy, the number of

delays with PUMAA is small when comparing to the fastest resource allocation (35.64 against 32.21). Moreover, in terms of delays, results are much better than the ones obtained by the uni-attribute auctions based upon price.

Discussion

This experiment shows that in a multi-attribute resource allocation environment all the attributes must be included in the decision processes. The quality of the allocations can be jeopardized otherwise (obtaining a good result regarding one attribute but a poor performance regarding the other ones). Thus, there is a need to reach a compromise between all attributes involved in the problem.

In this experiment we have seen how allocations based only upon price obtain the highest number of delays whilst ones which focus only on time provoke high costs. Multi-attribute methods achieve a compromise, as the number of delays produced by PUMAA is not much higher than the ones produced when picking the fastest resource. PUMAA obtained a similar price to the one obtained by the Vickrey auction allocation method.

6.3.2 Experiment 2: Strategy proofness

The goal of these experiments is to study the incentive compatibility of PUMAA. For that purpose we first studied if a resource agent would gain more from lying or bidding truthfully. In order to do so, we compared the profits (the difference between their costs and their incomes) which bidders obtained when all bid truthfully with the profits they obtained when some cheat. The experiment is first performed using the synthetic data simulation set (Scenario 1) and then using the production factory data set (Scenario 2). In these two scenarios, we expect that cheating agents should obtain fewer profits when cheating than when bidding truthfully; moreover we expect that this fact will not increase the auctioneers' expenses.

In a third scenario we wanted to verify if an intelligent agent which has no prior knowledge of the mechanism can learn that truthful bidding is the dominant strategy. For this purpose we introduce a learning agent into the market which uses reinforcement learning to identify the optimal bidding strategy. We expect that, as PUMAA is incentive compatible (see Chapter 3), truthful bidding will be the strategy with the highest probability of being selected.

Scenarios

The scenarios of this experiment are the following:

- **Scenario 1:** The main goal of this scenario is to evaluate whether bidder agent will profit more from lying or bidding truthfully. To test the hypothesis that truthful bidding is the best choice, the experiment is first conducted with two cheating agents (*RP1* and *RP2*) which provide lower delivery times values than the ones they plan to deliver. The experiment is conducted with seven honest agents (*RP3* to *RP9*). The experiment is repeated and the bidding strategy of *RP1* and *RP2* is modified from cheating to truthful bidding. In this scenario synthetic data is used to perform the simulation. There are three service agents (*SA1* to *SA3*) defining auction tasks in order to end their workflows whilst nine resource providers (*RP1* to *RP9*) compete for obtaining the auctioned tasks.

The allocations are performed using PUMAA with the product ($b * t$) defined as evaluation function. The length of the simulation is 200 time cycles (or days) and each simulation is repeated 200 times in order to obtain statistically significant data.

- **Scenario 2:** The goal of this scenario is to evaluate the strategy proofness of PUMAA in a scenario based upon real-data. Consequently, this scenario is similar to the previous one but uses production factory data for the simulation. The goal of this scenario is to evaluate the strategy proofness of PUMAA in a scenario based upon real-data. Consequently, this scenario is similar to the previous one but uses production factory data for the simulation. There are seven service agents *SA* (one for each type of task) which auction tasks. Eight resources providers *RP* (one for each type of provider) compete for obtaining these tasks. The scenario is executed two times: first with four agents providing untruthful bids (*RP5* to *RP8*) and then with all the agents offering truthful bids.

As in the previous scenario, the allocation is performed using PUMAA and the product as the evaluation function. The length of the simulation is also the same as the previous scenario (200 time cycles and 200 repetitions).

- **Scenario 3:** The goal of this scenario is to study which bidding strategy is chosen by an intelligent agent which has no prior knowledge of the market. To this end, we will use the synthetic data with 3 service agents and 9 resource providers. Eight resource providers (*RP2* to *RP9*) use an adaptive bidding strategy (they increase or decrease their economic bid depending on whether they win or lose in order to try to maximize




Figure 6.9: Resource providers' profits when *RP1* and *RP2* cheat (grey) and when they all bid truthfully (black).

their benefits). On the other hand, (*RP1*) tries to learn its optimal bidding strategy using a reinforcement learning algorithm [58].

Roughly, to learn the best strategy, *RP1* creates a table with all the strategies it can follow and assigns a probability of being chosen to each strategy. In the table, in addition to truthful bidding, other bidding strategies are considered so the learning bidder can test if they are profitable. Every time *RP1* participates in an auction, it updates the probability of choosing the the strategy used according to the utility it derived from it.

To simplify *RP1* learning, we define its bid range: it can bid from 0.5 to 2 times its true value (obtaining a total of 441 possible bidding strategies). For instance, if *RP1* provides a bid $B_{RP1} = (2b'_{RP1}, 0.5t'_{RP1})$ it means that *RP1* asks for the double of its true economic cost and that it says that lasting the task will take half the time of its actual duration.

The length of the simulation is 200 days and its repeated 200 times.

The results of the two first scenarios are measured in terms of bidder's profits and agents expenses (how much do they have to pay for allocating their tasks). In the third scenario we evaluate the probability of each strategy of being chosen, where the strategy with the highest probability will be the dominant strategy.




Figure 6.10: Service agents expenses when *RP1* and *RP2* cheat (grey) and when they all bid truthfully (black).

Results: Scenario 1

This scenario provides a first assessment of PUMAA's behavior in front of cheater agents. Figures 6.9 and 6.10 show the mean profits obtained by each bidder and the auctioneer expenses when all the agents bid truthfully (black bars) and when *RP1* and *RP2* use cheating as bidding strategy (grey bars).

The results of the first figure clearly show that *RP1* and *RP2* (dishonest) obtain more benefits when they are bidding according to their true values. Specifically *RP1* obtains 15.29% more profit when it bids truthfully than when it tries to cheat (8,810.38€ vs. 7,651.03€). *RP2* obtains 20.73% more profits (7,001.02€ vs. 5,798.56€). Throughout the rest of resource providers, we cannot find a clear trend in terms of profit variation. When cheaters bid truthfully, some of the honest agents see their benefits increased (*RA3*, *RA6* and *RA8*) whilst the others lose profits (*RA4*, *RA5*, *RA7* and *RA9*). However, this variation is never higher than a 5% (*RP9*).

If we assess the amount of money each auctioneer needed to allocate their tasks (Figure 6.10) it can be seen that there are no significant changes between their costs: when all agents bid honestly, *SA3* has to pay 2.10% less whilst *SA2* and *SA3* pay a 1.39% and a 1.38% more respectively. Student's t-test confirms that there are no difference between the




Figure 6.11: Mean resource providers' benefits when *RP5* to *RP8* cheat (grey) and when they all bid truthfully (black).

auctioneers' expense with a 99% of confidence.

Thus, this experiment corroborates that bidders obtain higher benefits by bidding truthfully in comparison to trying to deceive the auctioneers. It also shows that the strategy chosen by the bidders does not significantly affect the auctioneers' incomes due the fact that they pay less money when they obtain worse attributes than expected.

Results: Scenario 2

The goal of scenario 2 is to evaluate PUMAA's behavior when agents provide bids with non-truthful values in a real data-based simulation. Figure 6.11 shows the profit levels which agents obtain after performing the tasks they competed for. Figure 6.12 shows the mean expenses that the auctioneers had to invest in order to outsource their tasks and, finally, Figure 6.13 shows the utility which service agents obtain from outsourcing those tasks.

The first remarkable thing that can be observed in this experiment, is that all the cheating bidders (*RP5* to *RP8*) would have obtained higher benefits by bidding truthfully. Even the agents which do not obtain many allocations, due to their lack of skills (*RP6* and *RP8*) obtain higher profits if they reveal their true values. In this scenario the differences between the benefits which agents obtained bidding truthfully and the benefits they obtained when cheat-

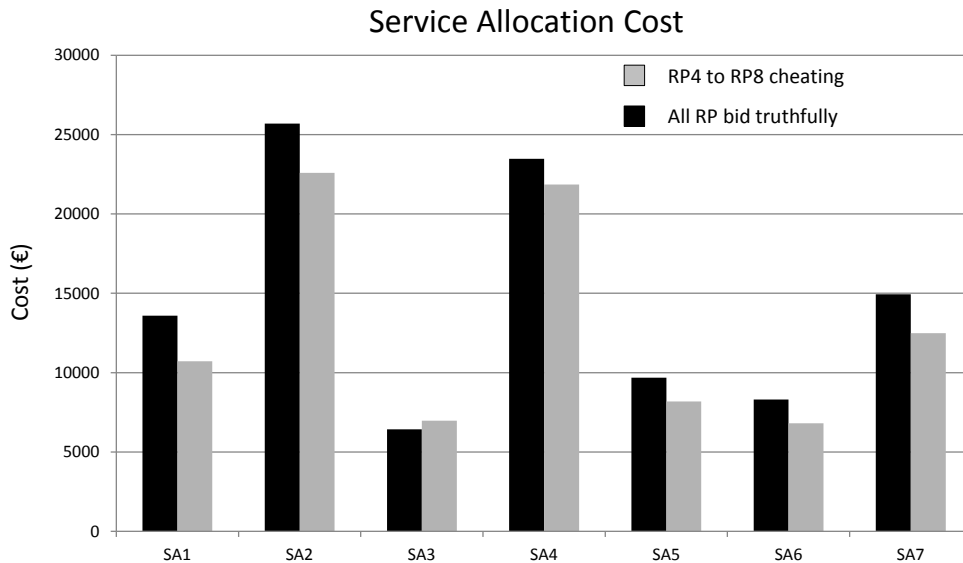


Figure 6.12: Mean service agents' expenses when *RP5* to *RP8* cheat (grey) and when they all bid truthfully (black).

ing are much more accentuated than in the previous example. *RP8* obtained four times more benefits by bidding truthfully, 900.26€ , than by cheating, 196.18€ ; *RP7* approximately double (14,728.48 vs. 7,454.84€) and *RP5* got 2.8 times more benefits by bidding truthfully (5,437.64 vs. 1,899.02€). On the other hand, bidders which bided truthfully in both cases, as in the previous scenario, obtained similar benefits (*RP2*, *RP3* and *RP4* up to a 7% more whilst *RP1* obtains an 8.02% less) despite the other agents strategies.

If we assess the price which each service agent has to pay in order to allocate all their tasks (Figure 6.12), we can see that the presence of cheating agents has decreased the cost of the allocations. This is brought about by the conditional payment of PUMAA: when a cheater agent takes longer than agreed to deliver a task its payment is reduced. In this scenario, as there is a high number of unreliable bids which are allocated a task, the agents with unreliable bids receive much less money than expected, provoking a reduction in the service agents expenditure. This fact also alters the bar chart of Figure 6.13, which presents the service agent's utilities. The chart shows that the utilities obtained when all the bidders bid truthfully and when some lie are very similar. If we perform a two sample Student's t-test, we can see that there is no difference between them with a confidence interval of 99%. When an agent delivers a task later than agreed, PUMAA reduces the payment in order to minimize the utility loss produced by this unexpected late delivery. When a resource provider fails to deliver the




Figure 6.13: Mean service agents' utilities benefits when *RP5 to RP8* cheat (grey) and when they all bid truthfully (black). Service agent utility is defined as $u_{SA}(T, p_i, dt'_i, dl_{SA}) = f_{SA}(T) - p_i * (\frac{dt'_i}{dl_{SA}})$ where $f_{SA}(T)$ is the valuation the agent gives to the task being performed, p_i the price it pays for it, dt'_i the time instant where it receives the task and dl_{SA} the task deadline

task as indicated in the bid, the service agent in charge of the task reduces the payment and preserves its utility.

Given that bidders obtain higher utilities by bidding truthfully than by lying, this scenario has been useful to point the incentive compatibility of PUMAA. However, the most interesting conclusion which can be drawn from this experiment is that PUMAA is able to keep the auctioneers' utility when bidders fail to deliver the items as agreed, showing its robustness.

Results: Scenario 3

The goal of the third Scenario is to test if an intelligent agent is able to learn its dominant strategy and, if so which one it picks.

Figure 6.14 illustrates the results of the smart agent's learning process, each bar corresponds to a possible strategy. The Z axis corresponds to the probability of choosing a certain strategy (between 0 and 1), the X axis corresponds to the delivery time attribute whilst the Y axis corresponds to the economic amount bid. On the X and Y axis a value below 1 corresponds to an overbidding strategy (saying that the agent will finish earlier than it can or asking for




Figure 6.14: Probability for each $RP1$ strategy after a 200 days simulation.

less money than its true value). A value higher than 1 corresponds to underbidding (offering values worse than true values), when both values are set to one the strategy corresponds to honest bidding (offering the agent's true values). The bar chart indicates that the strategy which has more chances of being chosen is honest bidding (at the very center of the chart, with a probability of 0.698), meaning that this is the strategy which has provided the highest utility to $RP1$. Therefore, after 200 days participating in the auction market, the agent has discovered that truthful bidding is its dominant strategy.

Figure 6.15 extends the previous information by showing the evolution of the strategy probabilities along time. The green line corresponds to the truthful bidding strategy whilst the clear red lines correspond to the rest. The plot shows that the honest bidding strategy is the strategy with a highest probability of being chosen. The chart also shows that at the initially the agent does not know which strategy to pick and that none of the strategies has much higher probabilities than the others, however, once the agent decides to try truthful bidding, the probability of this strategy increases fast until arrives to a probability of 0.7%. Given that there are 144 bidding strategies, the fact that the agent learns that that truthful bidding is the best strategy in less than 160 auctions is a remarkable fact. However, it must be considered




Figure 6.15: Evolution of the bidding strategy probabilities for the reinforcement learning agent.

that, as reinforcement learning is a non-deterministic algorithm, the agent may take longer to learn such information. In the plot it can also be observed that other strategies have increased their probability of being picked (despite being much smaller than the probabilities of truthful bidding). However, such strategies do not reach the same probabilities as honest bidding because at certain points they provide negative utility to bidders. As can be seen in Figure 6.14 these strategies correspond to the ones close to honest bidding (providing values near the true value).

The fact that an intelligent agent learns that truthful bidding is the best strategy does not prove the incentive compatibility of the mechanism as the agent may have been mistaken during the learning process. Moreover, the strategy learned may also have been biased by the strategies used by its competitors. However, if we take into account the demonstrations of Chapter 3, we can conclude that PUMAA is strategy proof.

Discussion

The three scenarios of this experimental section were targeted towards the corroboration of the claims in Section 3.4. All the experiments performed point to the incentive compatibility of PUMAA, supporting the statements of Chapter 3.

The first two scenarios demonstrated that a bidder obtains higher profits when bidding truthfully than when providing false bids. To test this finding we compared the benefits of a bidder

Resources	RP1	RP2	RP3	RP4	RP5	RP6	RP7	RP8
Skills	A,15,2,300	B,15,1,285	C,10,8,270	D,14,2,285	A,12,6,250	B,13,4,255	C,12,5,235	D,13,5,240
	C,14,8,200	D,14,2,290	A,14,2,265	B,10,1,290	C,13,4,245	D,15,4,245	A,13,5,250	B,13,5,245
	B,10,9,290	A,10,9,300	D,15,9,210	A,14,8,205	D,12,5,255	C,12,5,240	D,12,6,255	A,12,6,250

Table 6.6: Example of resources skills randomly generated per each agent: (Resource type, delivery time, tolerance to errors and cost)

lying with the benefits it would have obtained in the same auction environment if bidding truthfully. This test has been performed with synthetic and real-based data, in both cases the hypothesis that being honest is the best strategy for bidders has been confirmed. However, this finding was most striking using the production factory dataset where the differences were much higher than in the synthetic dataset. This is probably due to the big difference between resource providers in terms of skills. In addition, the second scenario also stated that when agents fail to deliver tasks as agreed the utility of the auctioneers is preserved due to PUMAA's payment rule. In the third scenario we utilized an intelligent agent with a reinforcement learning algorithm in order to allow the discovery of its best strategy by reinforcement learning. The results show that the agent discovered that truthful bidding was the best choice.

The three experiments described experimentally support the statement that PUMAA is an incentive compatible mechanism. Moreover, the second scenario shows that PUMAA can be useful in preserving the auctioneers utility in cases where bidders misdeliver tasks.

6.3.3 Experiment 3: Impact of the evaluation function in the auction allocation

In the previous experiments we used the product as the evaluation function. Now, in order to illustrate the importance of the evaluation function within PUMAA, we pretend to analyze the behavior of PUMAA when different evaluation functions are used. To this end, we add a new attribute to the synthetic data set and we perform different simulations using three different evaluation functions. We expect that each evaluation function will favor different kinds of allocations (equilibrated, polarized, etc.).

Scenario

In this scenario we add an extra attribute to the synthetic data: an error tolerance attribute ($e \in [0, 100]$) which corresponds to the percentage of error which a service agent can accept in a task. For instance, an agent can buy a set of steel frames with a length of 100 centimeters with a 0.01% tolerance, meaning that all the items of frames must have a length between 99.99 and 100.01 centimeters.




Figure 6.16: left) Task mean service time with different evaluation functions (WS: Weighted Sum, Prod: Product, E.n: Euclidean norm). center) Task mean error tolerance. right) Task mean cost.

The simulation involves 3 service agents and 8 resource providers which follow an adaptive strategy. Each resource provider agent has the skills to performing three different tasks, with different qualifications (see Table 6.6). As can be seen in Table 6.6, providers from *RP1* to *RP4* have unbalanced attributes (a good time, a good error margin or a good price) whilst resources *RP5* to *RP8* have balanced attributes (an average price, error margin and service time). These differences will be useful when analyzing the different behavior of the mechanisms when a different evaluation function is used.

The situation presented will be simulated with three different configurations. In each simulation, a different evaluation function is used by the service agents. The weighted sum (as we are considering 3 attributes, we need to consider 3 weights, particularly we use $\mu_0 = 0.5$ for the price, $\mu_1 = 0.3$ for the delivery time and $\mu_2 = 0.2$ for the error tolerance), the product and the Euclidean norm. Each scenario is repeated 200 times so statistically significant data can be extracted from the simulations.

We analyze the resulting allocations regarding the three attributes involved: time, error and price. We assume that bidders offer the Pareto set of attributes which maximizes their chances of winning the auction and maximizes their utilities.

Results

Figure 6.16 shows the outcomes of the different simulations from the service agents point of view (mean delivery time, mean error and mean price). Figure 6.17 shows the results from the resource providers point of view (focusing exclusively in the bidder's profits).

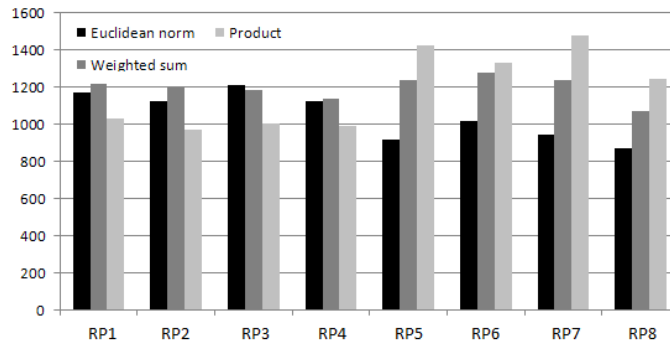


Figure 6.17: Resource providers benefits when using different evaluation functions.

In Figure 6.16 we can see that the different evaluation functions have changed the behavior of the auctions. When the weighted sums have been used the mean cost for each task is lower than when employing the other functions. If we compare the service time and the error margin we can see how the weighted sum provided worse results than the Euclidean norm and the product. This can be explained for the weights used in the evaluation function which gives more importance to the price in comparison to the other attributes. If we analyze the results of the product and the Euclidean norm we can observe that both evaluation functions have favored a higher task mean price. They favor a lower service time and error margin as none of the evaluation functions give priority to any attribute.

On the other hand, if we analyze the service providers benefits (Figure 6.17) we can see that the different evaluation functions have altered the bidder's incomes. For instance, we can see that the product (clear grey bars) favored the agents who provided a balanced value in all the attributes (*RP5* to *RP8*). Conversely, the euclidean norm has assigned more tasks to those bidders providing unbalanced attributes (*RP1* to *RP4*). This is explained by the nature of the Euclidean norm which gives lower values when the components of a vector are equilibrated, whilst it favors sets of attributes when one attribute has a very good value. Finally, we can see that upon utilizing the weighted sum there are not large differences between the agents. This is because the price attribute is the most weighted of the attributes and the allocations have been conditioned by the price bidders offered, not by the balance of their bids.

Discussion

This experiment points the importance of defining an appropriate evaluation function when performing an auction, as this clearly affects the resulting allocations. The auctioneer needs to use an evaluation function as similar as possible to its utility function, otherwise the resulting

allocations would not maximize its utility.

Particularly, in this experiment we compared the kind of allocations which the weighted sum, the product and the Euclidean norm would produce. We have observed that whilst the product favors allocations with balanced attributes, the Euclidean norm favors polarized bids with one good attribute. Regarding the weighted sum, the influence of each attribute will be conditioned by the weight the auctioneer gives to each.

Therefore, depending its the needs, the auctioneer will have to decide which method it will use. If the auctioneer needs a good balance amongst attributes the best method would be to use the product. If the auctioneer is interested in maximizing the value of one of the attributes without a particular preference it should use the Euclidean norm. Finally, if the auctioneer intends to focus upon a particular attribute (e.g. the price) the best method is to utilize the weighted sum and to assign a high weight to the desired attribute.

6.3.4 Experiment 4: Robustness to cheaters and utility preservation

The goal of this experiments is to study the robustness of the mechanism against cheating and to assess how auctioneers' utilities are affected by the presence of untruthful bidders. We compare our proposal (PUMAA) with the closest approach proposed in the literature: Che's second-price auction [15] (Che SP).

For this purpose we first execute a set of simulations with no dishonest agents, then we repeat the experiment replacing one honest bidder with one dishonest bidder. We repeat this process until all the agents within the simulation follow a cheating strategy. Finally we compare the number of delays produced within the system and the mean utilities which auctioneers have obtained at each of execution. This allows the study of variations caused by the number of dishonest agents.

In this experiment we expect that, without cheaters, PUMAA and Che's auction will provide very similar utilities given they have an analogous WDP. However, when cheaters appear we hypothesize that auctioneers will have higher utilities with PUMAA, due to its two-case payment mechanism. Regarding delays, we assume that both mechanisms will present similar numbers.

Scenario

This experiment is performed using the production factory dataset with 7 service agents (one of each service type) and 8 resource providers (one of each resource type). The agents consider two attributes: price and delivery time.

The experiment is executed 9 times (with 200 simulations with a length of 365 days for each execution). During the first execution all the agents follow a truthful bidding strategy. At the second execution one resource provider changes its strategy from truthful bidding to untruthful bidding (providing false delivery times and prices). This process is repeated until, at the last execution, where all the agents provide false delivery times. In all executions (both Che's second-price auction and PUMAA), service agents' utilities are defined as follows:

$$u_{SA}(T_{SA}, p_i, dt'_i) = v_o(T_{SA}) - p_i * \frac{dt'_i - st}{dl - st} \quad v_o(T_{SA}) = \begin{cases} \frac{1000}{emergency\ level(T_{SA})} & \text{if task on time} \\ 0 & \text{otherwise} \end{cases} \quad (6.1)$$

where T_{SA} is the task the service agent wants to allocate, dt'_i the delivery time which RP_i actually provided, st the time instant when the task could be started, dl the task deadline, p_i the payment that SA will perform, $v_o(T_{SA})$ the value which the SA gives to the task being perform and $emergency\ level(T_{SA})$ the emergency level of the task T_{SA} (note that the emergency level determines the task deadline and that the emergency level of each task is determined by the probability distributions specified in the dataset).

In PUMAA we use the following evaluation function (see Chapter 3):

$$V_{SA}(b_i, dt_i) = b_i * \frac{dt_i - st}{dl - st} \quad (6.2)$$

where b_i and dt_i are the economic cost and the delivery time which RP_i bid.

Whilst in Che SP we use the following score function and WDP:

$$S_{SA}(T_{SA}, b_i, dt_i) = v_o(T_{SA}) - b_i * \frac{dt_i - st}{dl - st} \quad (6.3)$$

The results of this experiment are presented in terms of utility, delays and payment per service.

Results

The goal of this experiment is to evaluate PUMAA's robustness to dishonest bidding agents and to compare it with Che's second price auction. Figure 6.18 shows the results obtained.




Figure 6.18: Comparison between Che's second price auction (Che SP) and PUMAA in terms of robustness. Top left: Mean total utility of the service agents in relation to the number of dishonest agents in the auction market. Top right: Normalized mean total utility of the service agents in relation to the number of dishonest agents in the auction market. Bottom left: Number of delays produced depending on the number of dishonest agents in the system. Bottom right: Mean payment performed by service agents depending on the number of dishonest agents in the system.

Top plots show the average total utility which service agents have obtained in the different simulations depending on the number of dishonest agents in the system; the left one shows this measure in absolute values whilst the right plot shows its normalized value in order to compare PUMAA and Che SP. The bottom left plot shows the mean number of delays produced in the system depending on the number of dishonest agents. Finally, bottom right plot shows the mean payment per service which service agents have made to resource providers.

Upon observing the utility which auctioneers obtained we can see that when using Che's SP, utilities rapidly decrease when more dishonest agents appear in the market. When all the bidders within the system bid dishonestly the auctioneers' utilities drop by up to 50%. The same phenomena occurs with PUMAA, but much slower than with Che's SP. We can see that when all bidders within the system act untruthfully the utility drops only to 78.92%. If we

analyze the utility values when half of the bidders still bid truthfully we can see that with PUMAA the auctioneers' utility has hardly changed (it preserves 95.74% of its original utility). Conversely, with Che SP, the auctioneers' utility decreased to 66.43% of the original utility. If we take into account that when all the bidders bid truthfully, Che SP and PUMAA displayed similar behavior and obtained similar utility for service agents.

If the number of delays which PUMAA and Che SP suffered when the number of cheating bidders increased is assessed, we can see that once again they have a very similar behavior. This is because PUMAA and Che SP have a very similar winner determination methodology (maximizing the expected utility function $u_0(B_i)$ is analogous to minimizing $V_0(B_i)$), thus, they tend to pick similar bids. This analogous behavior regarding the winner determination problems also results in similar utility when dishonest bidders are absent from the market. In an ideal world where all the participants were able to accurately estimate their skills and where every agent bid truthfully PUMAA and Che SP would result in the same outcomes.

The similarity between the delays produced by both mechanisms may seem inconsistent with the difference between the agents' utilities. However, if we observe Figure 6.18 bottom right, we can see that, whilst Che SP did not greatly reduce payments, as more dishonest agents appeared, PUMAA decreased the payment that service agents made to resource providers. This is due to the difference in the payment mechanisms of Che SP and PUMAA. Whilst Che SP attempts to match the second best bid score (no matter if the bidder respected its delivery time or not), PUMAA follows a two-case payment which preserves the auctioneer's utility in the case of the service agent not respecting the bidden attributes. In both cases the bidder receives a smaller economic amount than the one in asked for in the bid. However, in Che SP this reduction appears not to be enough to compensate the auctioneer utility loss when a delay occurs. This different payment methodology explains why when there is a high number of cheaters, when using Che SP agents lose much utility whilst with PUMAA they preserve it.

Discussion

This experiment evaluated the behavior of PUMAA when dealing with dishonest agents. The experiment showed that, unlike previous approaches such as Che's second price auction, when bidders fail to deliver a task or an item as agreed the utility of the auctioneers' is preserved.

The reason for this preservation of utility is the conditionality of PUMAA's payment mechanism. The experiment showed that despite the number of delays increasing when bidders acted untruthfully, the payments they received were reduced which meant that the auctioneers' utilities were almost preserved.

This experiment also stated that the auctioneers' utility cannot be completely protected given that when all the agents provided false delivery times the auctioneers' utility was slightly reduced. This situation is reasonable as, in certain circumstances, finishing a task (or delivering an item) after a deadline cannot be compensated by a payment reduction. In these cases we could say the utility is not preserved but loss is minimized.

6.4 Multi-dimensional Fairness Results

This section presents the experimentation and tests which illustrate the multi-dimensional fairness mechanisms presented in Chapter 5. To test multi-dimensional fairness approaches we use fair-PUMAA, the PUMAA extension built with FMAAC which allows the use of auctioneer-provided attributes to alter auction outcomes. Thus, the results of this section also correspond to the work described in Chapter 4.

To test fair-PUMAA and the multi-dimensional fairness mechanisms we present 3 different experiments. In the first one we compare uni-dimensional fairness with multi-dimensional fairness mechanisms. We also take advantage of this experiment to compare different priority methods presented in this thesis. The second experiment focuses upon the stochastic priority method; particularly upon the influence of priority in the wealth rank disorder. Finally, the third experiment studies how multi-dimensional fairness can be used to ameliorate the bidder drop problem.

In these experiments we use the production factory data set as it provides an environment where resource providers obtain differentiated outcomes due to their different skills for performing tasks. The evaluation function used to determine the winner of the auctions and its payment is the following:

$$V_0(B') = b_i * \frac{dt_i - st}{dl - st} * \frac{1}{1 + w_i} \quad (6.4)$$

where dt_i is the delivery time offered by RP_i , st the time instant when the task could be started, dl the task deadline, b_i the economic amount offered by RP_i and w_i the priority which the auctioneer gives to RP_i .

Consequently, the payment is computed as follows:

$$p = \begin{cases} \frac{(b_2 * (dt_2 - st)) * (1 + w_1)}{(dt_1 - st) * (1 + w_2)} & \text{if } dt' \leq dt_1 \\ \frac{(b_1 * (dt_1 - st))}{dt' - st} & \text{otherwise} \end{cases} \quad (6.5)$$

where the subindex (e.g. b_1) discerns between the best (1) and the second best bid (2), and dt' is the actual provided delivery time.

6.4.1 Experiment 1: Uni-dimensional vs. multi-dimensional fairness

The goal of this scenario is to investigate the differences between uni-dimensional fairness and the multi-dimensional fairness mechanisms presented in this thesis. In this experiment

we compare 6 different types of priority: not using a priority (which will act as a baseline), a uni-dimensional priority and the four deterministic priorities described in Chapter 5.

We expect that the uni-dimensional fairness method will provide the most egalitarian allocations in terms of cost but a bad performance in terms of delays. On the other hand, we expect that multi-dimensional fairness will solve this dichotomy obtaining good results both in terms of fairness and delays produced.

Scenario

In this scenario we use the production factory dataset. There are 7 service agents (one for each type of task) which outsource tasks to 8 resource providers (one for each type of resource). Resource providers follow an adaptive bidding strategy. We use 6 different configurations for fair-PUMAA: 1 using a uni-dimensional priority, 4 using multi-dimensional priorities and one without using priorities.

The uni-dimensional fairness method we use to compare with our proposal is based only upon the price, as for the previous mechanisms described in the literature [54]. Particularly, we define the priority in the uni-dimensional approach as follows:

$$w_i = 1 - \frac{\sum_{au=0}^{par(a_i)} r_i^{au}}{\sum_{l=1}^n \sum_{au=0}^{par(a_l)} r_l^{au}} \quad (6.6)$$

where w_i indicates the priority of the bidder a_i , $par(x)$ indicates the number of auctions in which an agent x has participated, r_i^{au} the revenue which bidder a_i obtained in the auction au and n the number of bidders in the market. This priority relates the economic amount a bidder has won compared to the total economic amount which has been paid by the auctioneer.

For multi-dimensional fairness methods, we test WOC, LS, BBWOC and BBLs, leaving stochastic approaches for another scenario. Concerning the parameters required by the priority methods, the values of ml in the LS and BBLs methods have been set to 70 (as in the simulation, on average, a resource provider participates in 10 auctions per day and we considered that a week without working is enough time to induce an agent to leave the market). To complement the experiment, a mechanism without fairness capacity (no priority) is also taken into account.

The simulations in this experiment last 400 cycles which correspond, to approximately a year of work in industry and each experiment is repeated 200 times in order to obtain statistically significant data.

This experiment will be evaluated in terms of price and delivery time (which are the attributes involved in the auction) and fairness:




Figure 6.19: Service agents' (auctioneers') mean expenditure at the end of simulations. No priority (dark blue), uni-dimensional priority (orange), LS (red), WOC (green), BBLs (purple), BBWOC (clear blue).

- Price: To evaluate how priority affects price, we compared revenue amongst bidders (the higher, the better) and the auctioneers' expenditure (the lower the better).
- Delivery time: as tasks are to be performed inside a specific time frame, we are interested in delays caused by tasks with an aim to avoid or minimize them. A one-dimensional fairness mechanism that does not take into account this attribute, but only the price, should fail to reduce delays. A multi-dimensional fairness mechanism similar to the one presented in this work, should achieve better results (minimize delays).
- Fairness: The Gini's index [31] is used to evaluate the fairness of the resulting allocations at the end of the simulation (the lower, the best). Gini's index is thought to analyze the social welfare differences among populations, thus, we have adapted it to our domain. The percentage of the population which corresponds to each resource provider is determined using the number of types of task which can perform each service agent. E.g. *RP2* has a population of the 18.42% as it can perform 7 tasks out of 38, *RP5* a population of 7.89% (3/38), etc.




Figure 6.20: Resource providers' (bidders') mean revenue at the end of simulations. From left to right: No priority (dark blue), uni-dimensional priority (orange), LS (red), WOC (green), BBLS (purple), BBWOC (clear blue).

Results

Figures 6.19 and 6.20 show the mean revenue and expenditure obtained by agents during the simulations with the different mechanisms considered in this scenario: no priorities, uni-dimensional fairness, and multi-dimensional fairness methods (WOC, LS, BBWOC and BBLS).

On the one hand, Figure 6.19 points that the use of priorities changes the costs of service allocation. It can be seen that, generally, when using priorities, service agents have not needed to increase their expenditure to allocate resources. The only case where an agent had to spend more money to allocate a task when using a priority was SA2. When using the LS method SA2 had to pay an average of 34,365.75 euros per simulation, when no priority was used it paid 31,652.89. In the rest of cases the agents' expenditure when using priority methods was lower than when no priority had been utilized. These results are as expected for a fairness mechanism which seeks to maintain the interest of agents in the auction. More agents means more competitive prices.

On the other hand, it can be seen that the use of priorities considerably changes the distribution of the revenues for resource providers (Figure 6.20): the differences between agents' revenues have been reduced as expected after the introduction of fairness mechanisms. An-




Figure 6.21: Resource agents' Gini's index for the resulting allocations in terms of revenue.

other remarkable finding is that bidders who obtained poor revenues (e.g. *RP5*, *6* and *8*) increased their benefits. It is important to note that the inclusion of fairness into the mechanism has modified the ranking of the richest service providers (e.g. *RP1* is the second richest agent when no fairness mechanism is applied whilst a bid-based fairness mechanism makes him the fourth richest agent). This result is analyzed further in the next experimental scenario.

To enhance the analysis of fairness within the allocations, Figure 6.21 presents a box plot of Gini's index corresponding to 200 executions of the scenario. It is important to note that the differences between agents revenues are modified by their skills. We observe that the four poorest agents (*RP5*, *RP7*, *RP8* and *RP6*) are the ones with least skills and it is reasonable to assume that other agents with better skills will be able to perform more tasks. This factor is taken into account in the fairness analysis performed: in the computation of Gini's index, each resource provider is weighted according to the number of skills (population) it is able to perform. The results obtained show that the use of priorities lowers the inequality coefficient. When no priority is used the index has an average value of 0.3623 whilst the use of priorities results in indexes lower than 0.3078 (with the BBLS). The lowest inequality index is obtained with the uni-dimensional priority (0.2134), followed by the BBWOC (0.2262) and the WOC (0.2371). The variance of the results does not allow us to determine which is the best priority method of those three (2-sample T tests do not find significant differences between the results). It can be seen that including a priority into the auction mechanism improves the equity of the




Figure 6.22: Mean number of delays produced for the different priority methods.

resulting allocation. As revenue only takes into account the first attribute of the bids, it is reasonable that the uni-dimensional priority obtained the best Gini's index. However, if we also take into account another attribute (time), results show a different face.

Figure 6.22 shows the mean number of delays produced during the simulations when using different types of priority. In this Figure, it is possible to observe that the uni-dimensional priority produces the highest mean number of delays, due to the fact that the uni-dimensional priority is only concerned with the cost attribute, pointing toward the need to apply fairness in a multi-dimensional way. When using multi-dimensional fairness, delays are significantly reduced, as shown in the box-plots. Using Student's t-test with a 99% of confidence we can confirm that multi-dimensional fairness mechanisms have overcome the uni-dimensional fairness approach. Inside the multi-dimensional priority methods, the LS is the worst and the WOC second worst. Regarding the WOC and the BB-WOC, the test does not observe significant differences between them. It can be seen that when using qualitative methods, the number of delays produced (7.242 and 8.003 time units) is similar to when no priority is used. This is due to the fact that bidders which offer almost-winning bids obtain high priority faster than low-quality bidders. In this way, weaker bidders still obtain advantage after a long period losing without jeopardizing the quality of the allocations.

In the boxplot of Figure 6.22, it is possible to observe that when no priority is used the number of delays produced is also low (7.911). This is due to the fact that the auction mech-

anism chooses the bids which minimize cost and execution time without taking into account priorities. However, the results obtained are worse than the BBWOC method (7.424), and very close to the BBL method (8.003). If Gini's index (Figure 6.17) is put together with delay results (Figure 6.18), we can conclude that multi-dimensional fairness mechanisms provide better allocation outcomes than no priority and uni-dimensional approaches in the long run. Of the qualitative methods BBWOC is the best.

Discussion

The experiment performed shows that the use of priorities reduces inequalities between resource providers. In this sense, uni-dimensional fairness presented the best results, followed by multi-dimensional priorities based upon WOC. The success of uni-dimensional priorities in terms of revenue equity can be explained by the fact that this priority only takes into account the agents' revenue. This focus on a single attribute had the effect of significantly increasing the number of delays in the system. In this sense, multi-dimensional fairness methods produced much better outcomes, as the number of delays they produced was very similar to the ones obtained when no priority was used.

Analyzing the results of this experiment, we can conclude that the use of multi-dimensional priorities, which act as an auctioneer provided attribute, can improve the fairness of the task allocation by assigning some services to the weaker agents. This difference in allocation has not increased the global cost of the allocation nor incurred significant delays. However the experiment also stated that by using priorities (no matter if they were uni or multi-dimensional), some of the strongest bidders may experience envy as their incomes could be reduced and they may lose position in the wealth ranking for the sake of equity.

6.4.2 Experiment 2: Stochastic priorities

The previous experiment showed that the use of priorities can modify the position which agents occupy in the wealth rank. As stated in Chapter 5, the use of stochastic priority methods can soften these modifications. The goal of this experiment is to analyze the role which probability plays in stochastic priority methods and how this influences the wealth ranking achieved by agents. We compare the stochastic probability methods when using a different update probability³ up_i (when $up_i = 0$ the priority is never updated whilst when $up_i = 1$ priorities are updated at each auction).

³Only auctioneer agents can have an update probability up .

It is reasonable to expect that results when up_i is 0.0 will be similar to the results of the previous scenario when no priority is used and that when $up_i = 1.0$ results will be similar to the ones obtained in the previous scenario with WOC, BBWOC, LS and BLS. With respect to the Wealth Rank Modification (WRM), ranking disorders are expected to increase as priorities are used more often (the higher the up_i , the higher WRM).

Scenario

In this experiment, we use the same scenario as in the first experiment (production factory dataset and 7 service agents outsourcing tasks to 8 resource providers) but using fair-PUMAA with the different stochastic priority methods (probWOC, probBBWOC, probLS and probBLS).

In order to test the influence of the priority upon the auction, this experiment is repeated several times with different priority update probabilities up_i , starting from $up_i = 0$ and increasing gradually by 0.1 until $up_i = 1$. All SA have the same up_i value. This process is performed with the four available stochastic priority methods.

The simulations in this experiment last 400 cycles and each execution (each priority method with its 11 possible update probabilities) is repeated 200 times in order to obtain significant data. The ml parameter of the probLS and probBLS methods is set to 70.

The results of this experiment will be evaluated in terms of fairness, using the Gini's index, and in how the wealth rank is modified (WRM).

With the WRM measure we want to evaluate the influence of fairness by comparing the position of an agent in the wealth rank when no fairness is used and when it is. For that purpose, we use the Spearman's footrule [24] distance (SFD).

In this way, we define the WRM as the SFD between the original wealth rank obtained when no priorities were used (R^o) and the wealth rank when fairness is used (R^f) divided by the SFD between the original wealth rank (R^o) and its reverse rank (R^r) which represents the maximum modification. Thus, the WRM measure is defined as follows:

$$WRM(R^f, R^o) = \frac{sfd(R^o, R^f)}{sfd(R^o, R^r)} \quad (6.7)$$

where $sfd(X, Y)$ is the Spearman's footrule between the ranks X and Y and $WRM \in [0, 1]$.

With this metric, a WRM of 0 means that the priority has not altered the wealth ranking at all whilst a WRM of 1 means that the wealth rank has been completely altered.




Figure 6.23: Mean Gini Index (blue circles) and mean WRM (green triangles) for the stochastic priority methods: probWOC (top left), probLS (top right), probBBWOC (bottom right) and probBBLS (bottom left).

For instance, if we have three agents a, b, c which obtained a wealth of 10, 9 and 3 respectively when no priority was used, and which obtained a wealth of 8, 9 and 5 when fairness is used we can see that there have been 2 displacements in the wealth rank (b goes from the second to the third whilst c varies from the third to the second). Given that inverting the original wealth rank produces 4 displacements (a varies 2 positions by going from the first position to the last and c 2 more positions from going from the last to the first), the WRM is 0.5 ($2/4$) since there has been 2 displacements out of 4 possible.

Results

Figure 6.23 shows the results obtained in terms of the Gini's index and the WRM with the probabilistic versions of WOC, BBWOC, LS and BBLs. As expected, in all cases, the resulting allocation is fairer when $up_i = 1$ as, with this probability, the method behaves in a deterministic

way. When analyzing the revenue of the agents, we can see how Gini's index progressively varies from an initial value (common to all the methods) situated around 0.36 when $up = 0.0$ to different values due to the priority calculation when $up = 1.0$. In all cases the mean Gini's index obtained when $up = 1$ (probWOC=0.2379, probBBWOC=0.2217, probLS=0.2688 and probBLS=0.3092) are similar to the values obtained in the previous scenario, confirming that the stochastic methods with $up = 1$ are equivalent to the deterministic approaches. If we assess the curve described by Gini's index, we can see that in probabilistic WOC and BBWOC the gradient of the curve increases as up_i rises. In the probabilistic LS and BLS the curve is flatter with a more stable gradient.

Figure 6.23 shows that fairness and rank preservation are related: the higher the fairness, the higher the disorder. In all cases, it can be seen that being fairer implies changing the wealth ranking, and this may result in complaints or dissatisfaction amongst bidders who lose position in the ranking. Consequently probBBWOC (the fairest mechanism) it is also the method with the highest WRM (0.3489) followed by probWOC (WRM=0.3092), probLS (0.2973) and probBLS (0.2901).

With respect to the compromise between fairness and rank preservation, it can be seen that the probWOC and probBBWOC methods with $up_i = 0.7$ may present the best solution as they obtain better fairness than probLS and probBLS, but with a lower rank modification. If we assess the line chart defining the WRM with probWOC and probBBWOC we can see that between $up_i = 0.4$ and $up_i = 0.7$ the curve is flatter than in the rest of the chart. Conversely, the curve corresponding to the Gini's index maintains a more regular slope for the up_i values. Therefore, we can conclude that $up_i = 0.7$ offers a good compromise point as there is not a big difference between $up_i = 0.4$ and $up_i = 0.7$ in terms of WRM but there is an important improvement in terms of Gini's index.

Discussion

The stochastic priority methods proposed herein offer the ability to reduce the WRM but not avoid it.

This experiment demonstrates that fairness and WRM are in conflict and that with fair-PUMAA one cannot be improved without compromising the other. Thus, not silver bullet can be provided and the analyzed methods must be configured according to the mechanism designer's needs. This can be achieved by tuning the up_i parameter which best suits to the problem. In this case, an update probability of 0.7 for probBBWOC seemed to be the best compromise between fairness and WRM.

6.4.3 Experiment 3: Tackling the bidder drop problem

The aim of this experiment is to evaluate how the use of priorities, defined in a multi-dimensional way, could minimize the bidder drop problem and other related problems. We assess how fair-PUMAA behaves when a group of bidders try to create an oligopoly working under production costs, forcing competitors to leave the market [35].

In this experiment a group of bidders starts offering bids with low economic costs (under true value) in order to obtain all the tasks in the market during a certain time period. With this strategy, a group of bidders intends to cause all competitors to leave so that they can create an oligopoly and control the market prices.

We evaluate the behavior of fair-PUMAA in these circumstances when no priority is used and when different types of multi-dimensional priorities are utilized. We expect that the use of multi-dimensional fairness will reduce the number of agents leaving the market and, consequently, price increases due to malicious oligopolistic practices will be avoided.

Scenario

This experiment uses the production factory data set, however, the number of agents involved in the market differs from the previous experiments. In this scenario there are 10 auctioneers which only auction tasks of type S5 (Table 6.4) and 50 resource providers which have the same skills (*RP6*) and true values (55€ per task, and the same execution times) but follow different bidding strategies. Half of them are honest and adaptive, whilst the other half are oligopolistic.

As the goal of this experiment is to analyze the bidder drop problem, bidders behavior is designed in such a way that if, after a certain number of auctions they have not obtained benefits, they leave the auction market and do not come back. On the one hand, twenty-five bidders follow an adaptive strategy improving their bids in order to win auctions [43]; they will remain in the market for up to 50 auctions without obtaining benefit. On the other hand, the oligopolistic agents begin following an underbidding strategy (working below their true value cost and incurring economic losses for a while). In this way, they try to force the adaptive agents to leave the market with goal for creating an oligopoly where they can fix the market price. As this happens, these agents try to rise prices so as to increase their benefits. In order to carryout this strategy, oligopolistic agents will remain for up to 100 auctions without obtaining benefits before leaving the market.

The experiment is performed with no priority set, WOC, BBWOC, probBBWOC ($up_i = 0.7$),




Figure 6.24: Active bidders in the market during a simulation where half of the agents (25) try to manipulate the auction working under their production cost and where the other half are adaptive.

LS and BBLs ($ml = 70$) as they have provided the best results in past experiments.

The results of this experiment are presented in terms of the task allocation, mean cost on behalf of the auctioneers (the closer to the bidders true values, the better) and the number of agents which remain in the system. This last element is used to evaluate the fairness and the bidder drop problem, in a fairer allocation bidders will preferentially stay on the market, thus a higher number of bidders reflects a fair allocation. The first metric can be used to analyze the asymmetric balance of negotiation power of bidders, whilst the second metric can be used to study the bidder drop problem. The results are also compared with an ideal situation where all the bidders always bid their true values and where all the agents stay in the market.

Results

Figure 6.24 shows the number of active bidders which are within the auction market along time when using different priority strategies and Figure 6.25 shows a boxplot of the mean active agents after the simulations. The first notable result is that priority methods maintain a higher number of agents interested in the market than a mechanism without priorities. In

the non-priority case, the chart shows that the number of participants drops drastically from 50 to 20.09, with all the adaptive agents thrown out (in this case all the adaptive bidders and even some oligopolistic bidders are forced to leave the auction due to the behavior of the most powerful oligopolistic agents). This situation is minimized when priorities are included in the auction mechanism. If we compare this result with an ideal scenario where all bidders act honestly and no bidder leaves the market, it means that only the 40% of the bidders are retained within the auction.

In Figure 6.24 it can be also observed that, no matter the priority method used, between the twentieth time cycle and the 70th, using any of the proposed priority methods, there is a reduction in the participation level. This phenomena is produced by the 25 agents which start underbidding in order to create an oligopoly. The influence of these bidders is lower depending on the kind of priority method used. The chart shows that the BBWOC (clear blue), the LS (black) and the WOC (grey) are the priority methods which best preserve the number of active bidders in the market (between 34 and 35 active bidders, which corresponds to the 68% and 70% of bidders of the ideal scenario). In these cases, priority has increased the adaptive agents chances of winning, keeping their interest in the auctions and avoiding the creation of an oligopoly.

The chart shows that the BBWOC (clear blue), the LS (black) and the WOC (grey) are the priority methods which best preserve the number of active bidders inside the auction (between 34 and 35 active bidders, which corresponds to the 68% and 70% of bidders of the ideal scenario). In these cases, priority has increased the chances of winning of the adaptive agents, keeping their interest in the auctions and avoiding the creation of an oligopoly by the side of the underbidding agents.

The probBBWOC priority calculation method and the BBLs priority do not appear to be as effective at preserving the number of active bidders: 30 (60%) and 25 (50%) respectively. In the first case this phenomena is probably produced by the reduced influence of the priority. In the second scenario, the lower number of active bidders is caused by the way priorities are assigned, based upon the bid fitness. As underbidding agents start offering *better* bids, in spite of their failure to win they obtain higher priorities than adaptive agents. In consequence, the chances of many adaptive agents winning an auction during the first part of the simulation decreases, and they leave the auction after a given losing streak.

Figure 6.26 states that the allocation costs are reduced when priorities are used. When no priority is used the mean allocation cost rises to an average of 72.65€ per task; however, when a priority method is used the mean cost is significantly reduced to 56.44€ (a 22.65%




Figure 6.25: Box plot showing the number of active bidders after the simulations when no priority is used and when using different methods of calculating the priority.

cheaper). Taking into account the data from this boxplot and the data in Figure 6.21 we can say that, under certain circumstances, the use of our priority methods can prevent and difficult bidder agents from fixing the market price.

Regarding the asymmetric balance of negotiation power, Figure 6.27 shows the evolution of the mean task costs during the simulation. The red line indicates the mean task cost, which a task would have had in the ideal case where all the bidders in the mechanism had bid truthfully (55€ per task). The first remarkable fact is that, independent of the priority method used, there is an initial drop in price. This is caused by the underbidding agents which desire to exclude the rest of bidders. After this initial gap it can be seen that the mean cost of the allocations increases progressively. However, the ceiling of this increase is different for each priority used. When no priority is used, as there are no adaptive bidders remaining in the market, oligopolistic agents can increase the price as much as they want (72.65€, a 33.09% higher than the true values) because they have more power to fix the market price. This difference is lower when priorities are used. BBWOC is the priority type which obtains the mean cost per task nearest to the one fixed by bidder's true value (56.93€ against 55€, only a 3.51% more expensive), very close to the ideal situation. The rest of the priority methods obtain higher mean price per task (between 57.97€ and 62.48€), the WOC resulted in the second lowest mean price and the BBLs method resulted in the worst. It is important to observe that the two methods which obtained the best performance are both based on the




Figure 6.26: Box plot showing the mean cost of a workflow allocation after the simulations when no priority is used and when using different methods of calculating the priority.

won auction coefficient. The supremacy of the WOC versus LS methods is due to the fact that losing streak methods favor agents who have not won in the last auctions, in consequence, when the underbidding agents have a large losing streak they may benefit from the priority. Conversely, the auction won coefficient takes into account all the victories agents have obtained over time. Thus, oligopolistic agents are ballasted as they have won almost all auctions at the beginning of the simulation, resulting in a lower chance of winning the rest of the auctions.

Discussion

This experiment has shown that the use of priorities can reduce the effects of the bidder drop problem, minimizing the effects of the unbalance of negotiation power between auctioneers and bidders. The fact that the number of bidders leaving the market is reduced also minimizes the asymmetric balance of negotiation power between bidders and auctioneers. This is because having a high number of active bidders prevents the formation of an oligopoly which fixes the price of the market (this occurs when no priority is used).

When assessing which is the best multi-dimensional priority for reducing the bidder drop problem, the experiment shows that BBWOC is the best option. Together with LS and WOC, BBWOC is the priority which maintains most bidders in the market; moreover, when it comes to the mean task cost, WOC favors prices near the bidders true-values.




Figure 6.27: Mean task cost after a simulation where half of the agents try to manipulate the auction by working under their production cost and where the other half adopt an adaptive strategy.

6.5 Summary

In this chapter we tested the methods proposed in Chapters 3, 4 and 5 using a multi-agent simulation environment which replicated the particularities of a supply chain domain. In the simulations, two data sets have been used: one synthetic and one built using real data. In both cases a set of service agents were acting as auctioneers in order to allocate their tasks to resource providers, who acted as bidders.

The first part of the chapter has been focused on evaluating PUMAA. Initially we studied how focusing upon a single-attribute during the allocation process compromises the quality of the results in terms of other attributes, thus pointing the need to use multi-attribute auctions. Experiments have supported the statements concerning the incentive compatibility of PUMAA described in Section 3.4. In the first set of experiments we saw how bidders obtained more benefits when they bid truthfully than when they tried to manipulate the mechanism. The experiment also showed that an intelligent agent without prior knowledge of the auction market, after a learning period, chooses honest bidding as the dominant strategy. The assessments also showed that the election of an appropriate evaluation function is a key factor for

the operation of PUMAA given that it will condition the resulting allocations and the balance among allocation attributes.

Regarding PUMAA's performance in comparison to other multi-attribute auction mechanisms, PUMAA overcomes other approaches such as Che's score auctions in terms of utility preservation. The experiments showed that, despite PUMAA not preventing the occurrence of unexpected delays, the conditional payment rule (which reduces the payment according to the submitted bid) can be used to preserve the auctioneer's utility in case of bidders not respecting the auction agreement.

In the second part of the chapter we assessed multidimensional fairness, our experiments showed that in multi-attribute auctions priority must be computed taking into account all the attributes involved in the auction. The quality of the allocation may be compromised otherwise. In our approach, being the winner or the loser of an auction means to taking into account all the attributes involved in the auction process. This was clearly observed when a uni-attribute priority was used: unlike the multi-dimensional priority methods which obtained fair allocations without increasing the number of delays, the uni-dimensional priority obtained slightly fairer allocations but jeopardized the quality of the allocations. When using uni-dimensional fairness the number of delays has increased by 6 times. Experimental results showed that improving the fairness of an allocation can alter the wealth distribution amongst agents and may annoy the richest agents. Stochastic methods can dampen this alteration without compromising to much the fairness of the allocation.

Finally, experiments showed that the use of a priority within the auction process can reduce the bidder drop problem and prevent the creation of oligopolies try to fix . This factor combined with the results of the first experiments making use of multi-dimensional priorities an interesting tool for mechanism designers who want to take into account equity in allocations in the long run in multi-attribute auctions.

Regarding mechanism design, it is important to highlight that fair-PUMAA has been developed using the FMAAC framework. FMAAC significantly reduces the difficulty of setting up new multi-attribute auctions by easily incorporating the attributes desired. fair-PUMAA has the same unverifiable and verifiable attributes as PUMAA but, thanks to the use of an auctioneer-provided attributes, it endows the mechanism with fairness properties.

CONCLUSIONS

This final chapter summarizes the goal, the methodology and the results presented in this thesis. It describes the contributions of this thesis to the research fields of multi-attribute resource allocation, mechanism design and multi-agent systems. Finally, it highlights possible future improvements and research directions which can be derived from this work.

7.1 Summary

The aim of this thesis was to provide a mechanism for producing multi-attribute resource allocations in uncertain and dynamic supply chain environments. Given the peculiarities of this domain, we were particularly interested in designing a multi-attribute auction mechanism for task and resource allocation. We wanted to consider economic costs and other relevant attributes for the supply chain domain (e.g. production times, qualities, etc.). Achieving satisfactory and efficient allocations using auctions in uncertain environments is a complex task as participants may be reluctant to reveal their real preferences (both in terms of cost and the rest of attributes involved in the allocation process).

With this in mind we designed PUMAA, a reverse Vickrey-based multi-attribute auction mechanism for allocating resources in workflow environments. PUMAA follows a reverse auction schema in which an auctioneer defines a task to be allocated and the set of attributes which will be used to evaluate the received bids and to determine the auction winner. The most relevant characteristic of PUMAA is its conditional payment mechanism. PUMAA conditions the payment rule on the basis of non-economic attributes offered in the winner bid. If the winning bidder delivers a set of attributes equal or better than the ones he offered, it receives a payment corresponding to the valuation of the second best bid (increasing the bidder's ex-

pected benefit). If the winner delivers worse attributes than the ones he offered, his payment is reduced in order to preserve the auctioneer's expected utility. This duality ensures that the mechanism is incentive compatible (due to its second price philosophy) but it also endows the mechanism with certain robustness. The payment reduction in case of not receiving as good attributes as expected allows the auctioneer to save part of their budget, which then can be used to amend this situation. They can then obtain better suppliers in future tasks.

Different aspects of PUMAA have been analyzed in this work. First, the incentive compatibility of the mechanism has been studied both from a theoretical and experimental standpoint, concluding that our proposal favors truthful bidding. Furthermore, we discussed and defined the requirements which an aggregation function must commit to in order to be used as an evaluation function in PUMAA and in other score-based auction mechanism; specifically, such an aggregation function should be strictly monotonic, real-valued and bijective within the attribute domain of the auction.

PUMAA has been tested using a multi-agent based simulation which emulates a supply chain domain. In this simulation, service agents try to outsource tasks whilst resource providers compete to perform them. Experimentation with PUMAA supported the hypothesis regarding the mechanism, corroborating its incentive compatibility. In addition, experimentation demonstrated that PUMAA can achieve a compromise between the different attributes involved in the allocation problems. Despite it not preventing the appearance of unexpected delays caused by untruthful auctions, it can preserve the auctioneer's utilities in such situations.

This thesis also analyzed the kinds of attributes involved in auctions. These can be classified as bidder-provided unverifiable attributes (those which act as the auction currency and whose value is only known by its owner), bidder-provided verifiable auctions (those attributes offered by bidders which its true value can be checked by the auctioneer at the end of the auction) and auctioneer-provided attributes (those which are introduced into the auction by the auctioneer itself and which express information and beliefs of the auctioneer regarding other agents).

Using the previous classification we developed FMAAC, a framework for multi-attribute auction customization based upon PUMAA. FMAAC adds auctioneer-provided attributes to PUMAA in order to allow the auctioneer to express information regarding the bidders it is interacting with. These attributes allow designers to customize the outcome of the auctions in order to meet their needs and to include new properties to the mechanism. For instance, we used FMAAC to create fair-PUMAA which includes a priority mechanism in order to produce egalitarian allocations instead of just utilitarian ones.

Fairness mechanisms are typically focused upon the incomes or benefits obtained by agents,

however, fairness mechanisms for multi-attribute resource allocation problems need to take into account all the attributes which affect the allocation as focusing only upon fairness concerning the participants' incomes may compromise the quality of the rest of the attributes involved in the decision process (e.g. delays in deliveries, alterations in energy consumptions, etc.). With this in mind, we proposed a multi-dimensional fairness mechanism for multi-attribute resource allocation. The proposed approach assigns higher priorities to the bidders who may leave the auction market due to bad results in previous auctions. As a result agents which have a high probability of leaving the market have a higher chance of winning an auction, thereby obtaining an incentive to remain in the market.

Fair-PUMAA and its multi-dimensional fairness mechanism were tested in the same multi-agent based simulator as PUMAA, using data extracted from a real manufacturing company. Experiments pointed that in multi-attribute domains fairness must be applied with all the attributes involved in the allocation taken into account, not just the economic cost. The results showed that whilst multi-dimensional fairness slightly outperformed uni-dimensional fairness, the application of uni-dimensional fairness compromised the quality of the allocations in terms of the rest of the attributes. The experiments also supported the hypothesis that multi-dimensional fairness can minimize the bidder drop problem.

7.2 Future Work

The methodologies presented in this thesis target the problem of allocating resources, in a dynamic way under uncertain production and supply chain scenarios. Despite completing the goals of this thesis, they also leave the door open to further research and new lines of investigation. This section discusses some ways in which the research conducted in this thesis can be continued.

PUMAA has proven to be a good method for the allocation of tasks in a supply chain environment, however, the approach we presented follows a sequential schema where the auctioneer allocates tasks one after another. Generalizing PUMAA for combinatorial auctions in order to allocate sets of tasks could be a challenging and interesting work [80]. An intuitive approach would be to follow the VCG auction schema which generalizes Vickrey auctions, however, it must be taken into account that PUMAA's dimensionality significantly increases the complexity of the winner determination process. Moreover, the introduction of combinatory allocations may compromise PUMAA's current properties and new studies should be performed.

Another interesting improvement for PUMAA could be the differentiation between agents

who lie on purpose and agents who accidentally under-perform. This differentiation could consist on the incorporation of a forgiveness notion (for instance, allowing an agent to occasionally fail to deliver its task and punishing it only if the failure is repeated along time) or with the implementation of a long-term payment taking into account not only the current task but also the tasks which an agent performed in past. This improvement, however, must be carried carefully as it would probably compromise the incentive compatibility of the mechanism.

Along the thesis we have assumed that the preferences of an auctioneer can be numerically represented. The use of similarity measures may allow skipping this assumption. For instance, the auctioneer could try to determine the winner of the auction by measuring how similar a bid is to the auctioneer's goal instead of evaluating the bids independently. Similarly, the auctioneer could compute the payment evaluating the similarity between the offered task and the delivered task. The use of a similarity function in the payment step may also allow the implementation of certain tolerance when determining if a bidder succeeded in performing a task. Nevertheless, this modification needs a deep study in order to evaluate which similarity measures can be used and to determine how they could affect the auction properties.

To exemplify FMAAC's behavior we presented a naive auction mechanism which included a trust parameter to PUMAA as an auctioneer-provided attribute (Section 4.2.6). Despite the fact that trust has not been explored in this thesis, a deep study of how to incorporate trust-methods (such as the ones described in [69, 70]) could be a particularly useful research for improving reliability in supply chain management. In the same line, extending the mechanism towards providing solutions in case of bid withdrawals would endow the mechanism with a stronger notion of robustness [12, 89].

Further research should also target a deep study of certain issues which may arise due to the inclusion of fairness into PUMAA. For instance, study should be made of how bidders which try to manipulate the market by using false identities may affect the mechanism, a common problem within fairness literature [54].

The research conducted in this thesis has focused on the supply chain domain, however, PUMAA's structure may be suitable for other domains dealing with multi-attribute resource allocation. We consider that PUMAA could be adapted and may be useful for other domains such as the electricity smart grid (e.g. matching demand and energy supply) or health care management (e.g. handling patient needs with the available resources in nearby health care centers).

Finally, recent research has identified that an agent's sense of utility is strongly conditioned by the context in which it is placed [38]. The same event may result in different satisfaction

levels depending on different contextual factors. Endowing PUMAA with new context-sensitive evaluation functions would increase the dynamism of the mechanism and could probably increase auctioneers satisfaction.

NOTATION GUIDE

This appendix summarizes the notation used in equations of Chapters 3, 4 and 5.

A.1 PUMAA

Notation used in Chapter 3:

- T^j **Task j**
 Task defined by m characteristics/parameters.
 For example, manufacturing 100 pens ($pa_1^j = 100$) of blue ink ($pa_2^j = blue$).

$$T^j = \langle pa_1^j, \dots, pa_m^j \rangle \quad (A.1)$$
- AR **Task attributes requirements.**
 It defines the conditions of the task service evaluated in the auction.
 For instance, delivery time (ar^1), quality (ar^1), etc.

$$AR = \langle ar^1, \dots, ar^n \rangle \quad (A.2)$$
- T_i^j **Agent a_i 's task j**
 Is the task j required by the agent a_i .
 When $i = 0$, it refers to the auctioneer.
- AR_i **Agent a_i 's task attributes requirements**
 Task attribute requirements required by the agent i .
 When $i = 0$, it refers to the auctioneer.

$$AR_i = ar_i^1, \dots, ar_i^n \quad (A.3)$$

- CFP **Call for proposals**

It contains the task to be auctioned and its requirements.

$$CFP = (T_0^j, AR_0)$$

$$CFP = (T_0^j, \langle ar_0^1, \dots, ar_0^n \rangle) \quad (A.4)$$

- B_i **Bid**

Bid of the agent a_i

It contains the economic cost b_i and the attributes offered by the agent a_i .

$$B_i = (b_i, AT_i) \quad (A.5)$$

- AT_i **Agent a_i 's offered attributes**

Set of attributes bided by agent i .

For instance an agent a_1 offers delivering the task in 60 minutes ($at_1^1 = 60$) and generating 10CO₂ kg ($at_1^2 = 10$).

AT_i refers to agent a_i bided attributes

AT_i' refers to agent a_i delivered attributes

AT_i^t refers to agent a_i true value attributes

$$AT_i = \langle at_i^1, \dots, at_i^n \rangle \quad (A.6)$$

Particularly AT_1 are the attributes of the winning bid and AT_2 the attributes of the second best bid.

- $v_i(T)$ **Value of a task T**

Valuation which an agent a_i gives to a task T .

Particularly, $v_0(T_0^j)$ corresponds to the valuation which an agent a_0 gives to a task T_0^j ; and $v_i(T_0^j, AT_i)$ corresponds to the valuation which a bidder a_i gives to a task T_0^j when providing a set of attributes AT_i .

- p_i **Payment**

Payment which an agent a_i receives after winning an auction

- u_i **Utility of an agent a_i**

It can correspond to an auctioneer (u_0):

$$u_0(T_0^j, p, AT_i) = v_0(T_0^j) - f_0(p, AT_i) \quad (A.7)$$

Or to a bidder ($i > 0$):

$$u_i(T_0^j, p, AT_i) = p - v_i(T_0^j, AT_i) \quad (A.8)$$

- \bar{u}_i **Expected utility of an agent a_i**

It can correspond to an auctioneer (\bar{u}_0):

$$\bar{u}_0(T_0^j, b_i, AT_i) = v_0(T_0^j) - f_0(b_i, AT_i) \quad (\text{A.9})$$

Or to a bidder ($i > 0$):

$$\bar{u}_i(T_0^j, b_i, AT_i) = b_i - v_i(T_0^j, AT_i) \quad (\text{A.10})$$

- $V_0(B_i)$ **Auctioneer's evaluation function**

Function of the auctioneer a_0 which rates a bid B_i .

Similar to Che's Score function S .

$$V_0(B_i) = V_0(b_i, AT_i) \quad (\text{A.11})$$

Note that V_0 will determine the winner of the auctioneer, thus, in order to maximize the auctioneer's utility, it must be as similar as possible to f_0 .

- $V_0^{-1}(x, AT_i)$ **Inverse function of the evaluation function**

Given a score x it returns the value of b_i .

$$V_0(b_i, AT_i) = x$$

$$V_0^{-1}(x, AT_i) = b_i \quad (\text{A.12})$$

A.2 FMAAC

Notation used in Chapter 4 which refers to FMAAC:

- A_i^p **Auctioneer-provided attributes.**

Attributes which the auctioneer add to the bid of agent a_i

For example a priority or a trust attribute.

D^p is the domain of the attribute.

$$A_i^p = (at_{1i}^p, \dots, at_{m_p i}^p) \text{ with } at_{1i}^p \in D_1^p, \dots, at_{m_p i}^p \in D_{m_p}^p \quad (\text{A.13})$$

- A_i^u **Bidder-provided unverifiable attributes.**

Unverifiable attributes of the agent a_i .

For example, the economic cost.

Acts as the auction currency.

We assume that $A_i^u = 1$.

D_i^u is the domain of the attribute.

$$A_i^u = at_i^u \text{ with } at_i^u \in D_1^u \quad (\text{A.14})$$

- A_i^v **Bidder-provided verifiable attributes.**

Unverifiable attributes of the agent a_i .

For example, delivery times or qualities.

D^v is the domain of the attributes.

$$A_i^v = (at_{1i}^v, \dots, at_{m_p i}^v) \text{ with } at_{1i}^v \in D_1^v, \dots, at_{m_p i}^v \in D_{m_p}^v \quad (\text{A.15})$$

- it **Item auctioned.**

Item it sold/bought in the auction.

In workflows usually it corresponds to a task (e.g. T_0^j)

- CFP **Call for proposals**

It contains the item to be auctioned and its attribute requirements.

It defines the bidder-provided attributes of the auction.

$$CFP = (it, AR_0)$$

$$CFP = (it, A_0^u, A_0^v) \quad (\text{A.16})$$

- v_i **Value of an item**

Valuation which an agent a_i gives to an item.

Particularly, $v_0(it)$ corresponds to the valuation which an auctioneer a_0 gives to an item; whilst $v_i(it, A_i^v)$ corresponds to the valuation which a bidder a_i gives to the item it when providing a set of verifiable attributes A_i^v .

- B_i **Bid**

It contains the bidder-provided attributes A_i^v and A_i^u by the agent a_i .

$$B_i = A_i^u \oplus A_i^v = (at_i^u, at_{1i}^v, \dots, at_{m_p i}^v) \quad (\text{A.17})$$

- B_i' **Extended bid**

It contains the attributes offered by the agent a_i plus A^p .

$$B_i' = A_i^u \oplus A_i^v \oplus A_i^p = (at_i^u, at_{1i}^v, \dots, at_{m_p i}^v, at_{1i}^p, \dots, at_{m_p i}^p) \quad (\text{A.18})$$

- u_i **Utility of an agent a_i**

It can correspond to an auctioneer (u_0):

$$u_0(it, p, AT_i^v) = v_0(it) - f_0(p, AT_i^v) \quad (\text{A.19})$$

Or to a bidder ($i > 0$):

$$u_i(it, p, AT_i^v) = p - v_i(it, AT_i^v) \quad (\text{A.20})$$

- \bar{u}_i **Expected utility of an agent a_i**

It can correspond to an auctioneer (u_0):

$$\bar{u}_0(it, B'_i) = v_0(it) - f_0(B'_i) \quad (\text{A.21})$$

Or to a bidder ($i > 0$):

$$\bar{u}_i(it, b_i, AT_i^v) = b_i - v_i(it, AT_i^v) \quad (\text{A.22})$$

- $V_0(B'_i)$ **Auctioneer's evaluation function**

Function of the auctioneer a_0 which assigns a value to a bid B_i .

Expressed in different forms.

$$\begin{aligned} & V_0(AT_i^u \oplus AT_i^v \oplus AT_i^p) \\ & V_0(AT_i^u, AT_i^v, AT_i^p) \\ & V_0\left(b_i, at_{1i}^v, \dots, at_{m_v i}^v, at_{1i}^p, \dots, at_{m_p i}^p\right) \end{aligned} \quad (\text{A.23})$$

Note that, to maximize the auctioneer a_0 's utility, V_0 must be coherent with f_0 .

- $V_0^{-1}(x, AT_i^v, AT_i^p)$ **Inverse function of the evaluation function**

Given a score x it returns the value of b_i .

$$V_0(b_i, AT_i^v, AT_i^p) = x$$

$$V_0^{-1}(x, AT_i^v, AT_i^p) = b_i \quad (\text{A.24})$$

- w_i **Priority attribute**

Priority attribute which acts as auctioneer-provided attribute.

- τ_i **Trust attribute**

Trust attribute which acts as auctioneer-provided attribute.

A.3 Multi-dimensional Fairness

Notation used in Chapter 5:

- w_i **Priority attribute**

Priority attribute which define the risk of a_i leaving the auction market.

$$w_i \in [0, 1]$$

The higher w_i , the higher the chances of a_i leaving the market.

- $won(a_i)$ **Auctions won**

Number of auctions which agent a_i has won.

- $par(a_i)$ **Auctions participated**
Number of auctions in which agent a_i has participated.

- $ls(a_i)$ **Streak**
Number of auctions which agent a_i has consecutively lost.

- ml **Maximum lost auctions threshold**
A bidder who has lost ml auctions is as susceptible to leave the market as one which has lost $ml + 1$ auctions

- q_i **Bid fitness**
Indicates how good was a bid B_i compared to the winning bid.
 $w_i \in [0, 1]$
A value of 1 means that B_i was as good as the auction winner.

- up_i **Update probability**
It is the probability of an auctioneer a_i of updating the bidders priorities.
 $up_i \in [0, 1]$
A value of 1 means that the priority is always be updated.
Conversely, A value of 0 means that the priority will never be updated.

Bibliography

- [1] K. Alexopoulos, S. Makris, V. Xanthakis, and G. Chryssolouris. A web-services oriented workflow management system for integrated digital production engineering. *CIRP Journal of Manufacturing Science and Technology*, 4(3):290 – 295, 2011.
- [2] S. Athey and G. Ellison. Position auctions with consumer search. *Quarterly Journal of Economics*, 126(3):1213–1270, 2011.
- [3] L. Ausubel and P. Milgrom. The lovely but lonely vickrey auction. In *Combinatorial auctions*, pages 17–40. The MIT Press, 2006.
- [4] C. Barrett, L. de Moura, and A. Stump. Smt-comp: Satisfiability modulo theories competition. In *Computer Aided Verification*, pages 20–23. Springer, 2005.
- [5] C. W. Barrett, R. Sebastiani, S. A. Seshia, and C. Tinelli. Satisfiability modulo theories. In *Handbook of satisfiability*, pages 825–885. IOS Press, 2009.
- [6] Y. Bartal, R. Gonen, and N. Nisan. Incentive compatible multi unit combinatorial auctions. In *Proceedings of the 9th conference on Theoretical aspects of rationality and knowledge (TARK03)*, pages 72–87. ACM, 2003.
- [7] M.-J. Bellosta, S. Kornman, and D. Vanderpooten. A unified framework for multiple criteria auction mechanisms. *Web Intelligence and Agent Systems*, 6(4):401–419, 2008.
- [8] M.-J. Bellosta, S. Kornman, and D. Vanderpooten. Preference-based english reverse auctions. *Artificial Intelligence*, 175(7):1449 – 1467, 2011.
- [9] M. Bichler. An experimental analysis of multi-attribute auctions. *Decision Support Systems*, 29(3):249–268, 2000.
- [10] M. Bichler, M. Kaukal, and A. Segev. Multi-attribute auctions for electronic procurement. In *Proceedings of the First IBM IAC Workshop on Internet Based Negotiation Technologies*, Yorktown Heights, NY, March 1999.

-
- [11] A. Blume, P. Heidhues, J. Lafky, J. Münster, and M. Zhang. All equilibria of the multi-unit vickrey auction. *Games and Economic Behavior*, 66(2):729–741, 2009.
- [12] M. Bofill, D. Busquets, V. Muñoz, and M. Villaret. Reformulation based maxsat robustness. *Constraints*, 18(2):202–235, 2013.
- [13] C. D. Cantrell. *Modern mathematical methods for physicists and engineers*. Cambridge University Press, 2000.
- [14] C. R. Carter, L. Kaufmann, S. Beall, P. L. Carter, T. E. Hendrick, and K. J. Petersen. Reverse auctions: grounded theory from the buyer and supplier perspective. *Transportation Research Part E: Logistics and Transportation Review*, 40(3):229–254, 2004.
- [15] Y.-K. Che. Design competition through multidimensional auctions. *The RAND Journal of Economics*, 24(4):668–680, 1993.
- [16] Y. Chevaleyre, P. E. Dunne, U. Endriss, J. Lang, M. Lemaitre, N. Maudet, J. Padget, S. Phelps, J. A. Rodriguez-Aguilar, and P. Sousa. Issues in multiagent resource allocation. *Informatica*, 30(1):3–31, 2006.
- [17] E. H. Clarke. Multipart pricing of public goods: An example. *Public Prices for Public Products*. Urban Inst., Washington, 1972.
- [18] S. Clauß and M. Köhntopp. Identity management and its support of multilateral security. *Computer Networks*, 37(2):205–219, 2001.
- [19] J. Collins and M. Gini. Magnet: a multi-agent system using auctions with temporal and precedence constraints. In *Multiagent based Supply Chain Management*, pages 273–314. Springer, 2006.
- [20] R. K. Dash, N. R. Jennings, and D. C. Parkes. Computational-mechanism design: A call to arms. *Intelligent Systems, IEEE*, 18(6):40–47, 2003.
- [21] E. David, R. Azoulay-Schwartz, and S. Kraus. Protocols and strategies for automated multi-attribute auctions. In *Proceedings of the first international joint conference on Autonomous agents and multiagent systems: part 1, AAMAS '02*, pages 77–85, New York, NY, USA, 2002. ACM.
- [22] E. David, R. Azoulay-Schwartz, and S. Kraus. Bidding in sealed-bid and english multi-attribute auctions. *Decision Support Systems*, 42(2):527–556, 2006.

- [23] L. De Moura and N. Bjørner. Z3: An efficient smt solver. *Tools and Algorithms for the Construction and Analysis of Systems*, 4963/2008:337–340, 2008.
- [24] P. Diaconis and R. L. Graham. Spearman’s footrule as a measure of disarray. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(2):262–268, 1977.
- [25] B. Edelman, M. Ostrovsky, and M. Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review*, 97(1):242–259, 2007.
- [26] U. Endriss, N. Maudet, F. Sadri, and F. Toni. Resource allocation in egalitarian agent societies. In *Secondes Journées Francophones sur les Modèles Formels d’Interaction (MFI-2003)*, pages 101–110. Cépaduès-Éditions, May 2003.
- [27] U. Endriss, N. Maudet, F. Sadri, and F. Toni. Negotiating socially optimal allocations of resources. *Journal of Artificial Intelligence Research (JAIR)*, 25:315–348, 2006.
- [28] H. Fu, P. Jordan, M. Mahdian, U. Nadav, I. Talgam-Cohen, and S. Vassilvitskii. Ad auctions with data. In *Algorithmic Game Theory*, pages 168–179. Springer, 2012.
- [29] J. Gao, Y. Yao, V. Zhu, L. Sun, and L. Lin. Service-oriented manufacturing: a new product pattern and manufacturing paradigm. *Journal of Intelligent Manufacturing*, 22(3):435–446, 2011. 10.1007/s10845-009-0301-y.
- [30] R. J. Garrat. Vickrey-groves-clark mechanism (economics university lecture). *University of California Santa Barbara*, 2012.
- [31] C. Gini. Variabilità e mutuabilità. contributo allo studio delle distribuzioni e delle relazioni statistiche (variability and mutability. contribution to the study of the statistic distributions and relations.). *Tipogr. di Cupini*, 1912.
- [32] M. Grabisch, J.-L. Marichal, R. Mesiar, and E. Pap. Aggregation functions. In *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, New York, NY, USA, 1st edition, 2009.
- [33] T. Groves. Incentives in teams. *Econometrica: Journal of the Econometric Society*, pages 617–631, 1973.
- [34] B. P. Harrenstein, M. M. de Weerd, and V. Conitzer. A qualitative vickrey auction. In *Proceedings of the 10th ACM conference on Electronic commerce, EC ’09*, pages 197–206, New York, NY, USA, 2009. ACM.

- [35] B. F. Hobbs. Network models of spatial oligopoly with an application to deregulation of electricity generation. *Operations Research*, 34(3):395–409, 1986.
- [36] K. Hoesch-Klohe, A. Ghose, and L.-S. Lê. Towards green business process management. In *IEEE SCC*, pages 386–393. IEEE Computer Society, 2010.
- [37] A. Holland and B. O’Sullivan. Robust solutions for combinatorial auctions. In *Proceedings of the 6th ACM conference on Electronic commerce*, pages 183–192. ACM, 2005.
- [38] J. S.-C. Hsu. Understanding the role of satisfaction in the formation of perceived switching value. *Decision Support Systems*, (0):–, 2013.
- [39] M. Jofre-Bonet and M. Pesendorfer. Optimal sequential auctions. Working paper, London School of Economics and Political Science, 2006.
- [40] S. Kraus. Negotiation and cooperation in multi-agent environments. *Artificial Intelligence*, 94:79–97, 1997.
- [41] V. Krishna. *Auction Theory*. Academic Press/Elsevier, 2009.
- [42] V. Krishna and M. Perry. *Efficient mechanism design*. Hebrew University of Jerusalem, Center for Rationality and Interactive Decision Theory, 1997.
- [43] J. Lee and B. Szymanski. A novel auction mechanism for selling time-sensitive e-services. In *IEEE Conference on ECommerce Technology (CEC’05)*, pages 75–82. Press, 2005.
- [44] J.-S. Lee. *Recurrent auctions in e-commerce*. ProQuest, 2007.
- [45] D. Lehman, R. Muller, and T. Sandholm. *The Winner Determination Problem (Chap. 12)*. MIT Press, 2006.
- [46] D. Lehmann, L. I. O’callaghan, and Y. Shoham. Truth revelation in approximately efficient combinatorial auctions. *J. ACM*, 49:577–602, September 2002.
- [47] M. Lemaître, G. Verfaillie, H. Fargier, J. Lang, N. Bataille, and J. m. Lachiver. Equitable allocation of earth observing satellites resources. In *Proceedings of 5th ONERA-DLR Aerospace Symposium (ODAS03)*, 2003.
- [48] J. Liker. *The Toyota Way : 14 Management Principles from the World’s Greatest Manufacturer*. Mcgraw-hill, 2003.

- [49] A. Likhodedov and T. Sandholm. Auction mechanism for optimally trading off revenue and efficiency. In *Proceedings of the 4th ACM conference on Electronic commerce*, pages 212–213, San Diego, CA, USA, 2003. ACM.
- [50] T. Mahr and M. M. de Weerd. Multi-attribute vickrey auctions when utility functions are unknown. In P.-Y. Schobbens, W. Vanhoof, and G. Schwanen, editors, *Proceedings of the Belgium-Dutch Conference on Artificial Intelligence (BNAIC)*, pages 221–227, Enschede, Netherlands, 2006. BNVKI.
- [51] V. Mañoso, F. J. Labrador, and A. Fernández-Alba. Tipo de distorsiones cognitivas durante el juego en jugadores patológicos y no jugadores (in spanish, cognitive distortion types during gambling in pathological players and unusual players). *Psicothema*, 16(4):576–581, 2004.
- [52] E. Maskin and J. Riley. Equilibrium in sealed high bid auctions. *The Review of Economic Studies*, 67(3):pp. 439–454, 2000.
- [53] R. P McAfee and J. McMillan. Auctions and bidding. *Journal of economic literature*, 25(2):699–738, 1987.
- [54] J. Murillo, B. López, V. Muñoz, and D. Busquets. Fairness in recurrent auctions with competing markets and supply fluctuations. *Computational Intelligence*, 28(1):24–50, 2012.
- [55] J. Murillo, V. Munoz, B. Lopez, and D. Busquets. A fair mechanism for recurrent multi-unit auctions. In *Multiagent System Technologies*, volume 5244 of *Lecture Notes in Computer Science*, pages 147–158. Springer Berlin Heidelberg, 2008.
- [56] J. Murillo Espinar. *Egalitarian behaviour in multi unit combinatorial auctions*. PhD thesis, Universitat de Girona, 2010.
- [57] R. B. Myerson. Incentive Compatibility and the Bargaining Problem. *Econometrica*, 47(1):61–74, 1979.
- [58] V. Nanduri and T. K. Das. A reinforcement learning model to assess market power under auction-based energy pricing. *Power Systems, IEEE Transactions on*, 22(1):85–95, 2007.
- [59] J. V. Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.
- [60] D. Parkes and J. Kalagnanam. Iterative multiattribute vickrey auctions. Technical report, Harvard University, 2002.

- [61] D. C. Parkes and J. Kalagnanam. Models for iterative multiattribute procurement auctions. *Management Science*, 51(3):435–451, 2005.
- [62] S. Parsons, J. A. Rodriguez-Aguilar, and M. Klein. Auctions and bidding: A guide for computer scientists. *ACM Computing Surveys (CSUR)*, 43(2):10:1–10:59, Feb. 2011.
- [63] A. Pla, P. Gay, J. Meléndez, and B. López. Petri net-based process monitoring: a workflow management system for process modelling and monitoring. *Journal of Intelligent Manufacturing*, pages 1–16, 2012.
- [64] A. Pla, B. López, and j. Murillo. How to demonstrate incentive compatibility in multiattribute auctions. In *Proceedings of the 13th International Conference of the Catalan Association for Artificial Intelligence*, volume 256 of *Frontiers in Artificial Intelligence and Applications*, pages 303–306. IOS Press, 2013.
- [65] R. Porter, A. Ronen, Y. Shoham, and M. Tennenholtz. Fault tolerant mechanism design. *Artificial Intelligence*, 172(15):1783–1799, 2008.
- [66] R. Prodan, M. Wiecek, and H. M. Fard. Double auction-based scheduling of scientific applications in distributed grid and cloud environments. *Journal of Grid Computing*, 9(4):531–548, 2011.
- [67] J. Puustjärvi. Coordinating electronic auctions by workflow systems. In *Proceedings of the IADIS International Conference*, pages 139–142, 2005.
- [68] R. I. Rackson, J. A. Krane, and P. Trevisani. Multiple auction coordination method and system (us6415270 u.s. patent), 07 2002.
- [69] S. D. Ramchurn, C. Mezzetti, A. Giovannucci, J. A. Rodriguez-Aguilar, R. K. Dash, and N. R. Jennings. Trust-based mechanisms for robust and efficient task allocation in the presence of execution uncertainty. *Journal of Artificial Intelligence Research (JAIR)*, 35:119–159, June 2009.
- [70] J. Sabater and C. Sierra. Review on computational trust and reputation models. *Artificial Intelligence Review*, 24(1):33–60, 2005.
- [71] N. Sadeh and J. Sun. Multi-attribute supply chain negotiation: coordinating reverse auctions subject to finite capacity considerations. In *Proceedings of the 5th international conference on Electronic commerce, ICEC '03*, pages 53–60, New York, NY, USA, 2003. ACM.

- [72] T. Sandholm and S. Suri. Improved algorithms for optimal winner determination in combinatorial auctions and generalizations. In *World Congress of the Game Theory Society*, pages 90–97. AAAI/IAAI, 2000.
- [73] R. Shah and P. Ward. Lean manufacturing: context, practice bundles, and performance. *op. Management*, 21(2):129–149, 2003.
- [74] D.-H. Shih, B. Lin, and S.-Y. Huang. Morvam: A reverse vickrey auction system for mobile commerce. *Expert Systems with Applications*, 32(4):1113 – 1123, 2007.
- [75] Y. D. Smet. Multi-criteria auctions without full comparability of bids. *European Journal of Operational Research*, 177(3):1433 – 1452, 2007.
- [76] K. Srnka and S. Koeszegi. From words to numbers: how to transform qualitative data into meaningful quantitative results. *Schmalenbach Business Review*, 59(1):29–57, 2007.
- [77] S. Strecker. *Multiattribute Auctions in Electronic Procurement: Theory and Experiment*. Dissertation, Universität Karlsruhe (TH), 2004.
- [78] S. Talluri and G. L. Ragatz. Multi-attribute reverse auctions in b2b exchanges: A framework for design and implementation. *Journal of Supply Chain Management*, 40(1):52–60, 2004.
- [79] M. Tavana, A.-R. Abtahi, and K. Khalili-Damghani. A new multi-objective multi-mode model for solving preemptive time–cost–quality trade-off project scheduling problems. *Expert Systems with Applications*, 41(4), 2014.
- [80] F. Torrent, A. Pla, and B. L’opez. Using multi-attribute combinatorial auctions for sustainable workflow scheduling. In *AAAI Artificial Intelligence Conference 2014*, page Submitted, 2014.
- [81] University of Girona. Production factory data set http://eia.udg.es/~apla/fac_data/, april 2013.
- [82] W. M. P. Van der Aalst. The application of Petri nets to workflow management. *The Journal of Circuits, Systems and Computers*, 8(1):21–66, 1998.
- [83] H. R. Varian. Economic mechanism design for computerized agents. In *Proceedings of the 1st USENIX Workshop on Electronic Commerce*, pages 13–21, Berkeley, CA, USA, 1995. USENIX.

-
- [84] H. R. Varian. Position auctions. *International Journal of Industrial Organization*, 25(6):1163 – 1178, 2007.
- [85] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, 16(1):8–37, 1961.
- [86] M. Wang and H. Wang. Intelligent agent supported flexible workflow monitoring system. In *Proceedings of the 14th International Conference in Advanced Information Systems Engineering (CaiSE2002)*. Citeseer, 2002.
- [87] WFMC. *Workflow Management Coalition Terminology and Glossary (WFMCTC-1011)*. Workflow Management Coalition, 1999.
- [88] M. Wooldridge. *An Introduction to Multiagent Systems*. Wiley, Chichester, UK, 2nd edition, 2009.
- [89] Y. Zhang and S. Verwer. Mechanism for robust procurements. In I. Rahwan, W. Wobcke, S. Sen, and T. Sugawara, editors, *PRIMA 2012: Principles and Practice of Multi-Agent Systems*, volume 7455 of *Lecture Notes in Computer Science*, pages 77–91. Springer Berlin Heidelberg, 2012.
- [90] D. Zhao, D. Zhang, and L. Perrussel. Mechanism design for double auctions with temporal constraints. In *Proceedings of the Twenty-Second international joint conference on Artificial Intelligence*, volume 1, pages 472–477, Barcelona, 2011. AAAI Press.