Auction-Based Allocation and Pricing of Telecommunication Resources

Marina E. Bitsaki

Doctoral Dissertation

Department of Computer Science
University of Crete
November 2004

ΠΑΝΕΠΙΣΤΗΜΙΟ ΚΡΗΤΗΣ
ΣΧΟΛΗ ΘΕΤΙΚΩΝ ΕΠΙΣΤΗΜΩΝ
ΤΜΗΜΑ ΕΠΙΣΤΗΜΗΣ ΥΠΟΛΟΓΙΣΤΩΝ

Εκχώρηση και Χρέωση Τηλεπικοινωνιακών Πόρων μέσω Δημοπρασιών

Διατριβή που υποβλήθηκε από την

Μαρίνα Ε. Μπιτσάκη

ως μερική απαίτηση για την απόκτηση του Διδακτορικού Διπλώματος

Ηράκλειο, Νοέμβριος 2004

Συγγραφέας:

Εξεταστική Επιτροπή:

Χρήστος Νικολάου, Καθηγητής, Πανεπιστήμιο Κρήτης, Επιβλέπων

Γεώργιος Δ. Σταμούλης, Αναπληρωτής Καθηγητής, Οικονομικό Πανεπ. Αθηνών, Συνεπιβλέπων

Απόστολος Τραγανίτης, Καθηγητής, Πανεπιστήμιο Κρήτης

Κώστας Κουρκουμπέτης, Καθηγητής, Οικονομικό Πανεπιστήμιο Αθηνών

Γρηγόριος Αντωνίου, Καθηγητής, Πανεπιστήμιο Κρήτης

Βασίλειος Σύρης, Επίκουρος Καθηγητής, Πανεπιστήμιο Κρήτης

Γεώργιος Φ. Γεωργακόπουλος, Επίκουρος Καθηγητής, Πανεπιστήμιο Κρήτης

Δεκτή:

Δημήτριος Πλεξουσάκης, Αναπληρωτής Καθηγητής, Πρόεδρος Επιτροπής Μεταπτυχιακών Σπουδών

Acknowledgements

I am very grateful to Professor George D. Stamoulis, for the effort and time he spent to guide me throughout my Ph.D. studies. He shared his knowledge with me and challenged me to go deep into research in the area of charging communication networks. His experience and research ability, going beyond ordinary standards, have been valuable to me. His tireless guidance, help and patience were instrumental in making this work a reality.

I would like to thank Professor Costas Courcoubetis who has contributed a great deal to this dissertation. I want to express my gratitude to him for giving me the opportunity to do research initially at the University of Crete and for the last years of my studies at Athens University of Economics and Business. His encouragement and support have helped me overcome the difficulties when it was most required.

I would like to thank Professor Christos Nikolaou for supporting my effort throughout my time at the University of Crete and for his confidence in me. I am greatly indebted to him for giving me the opportunity to continue and complete my studies at the University of Crete.

I would also like to thank Manos Dramitinos for the fruitful discussions and the excellent collaboration we had during my stay in Athens. His assistance in technical matters as well as his detailed comments on parts of this dissertation were valuable to me.

I warmly thank Professor George Polyzos for providing me access to the resources of the Mobile Multimedia Laboratory at Athens University of Economics and Business. Finally, I would like to thank all the members of the Networks Economics and Services group at Athens University of Economics and Business for their encouragement and the interesting discussions we had.

To Professor Ravi Mazumdar

Auction-Based Allocation and Pricing of Telecommunication Resources

Marina E. Bitsaki

Doctoral Dissertation

Department of Computer Science

University of Crete

Abstract

The growth of communication markets and the variety of services being offered has encouraged the study of pricing issues in communication networks. Besides the need for charging those using the network premises, pricing plays a major role in controlling congestion and induces incentives to utilize the network in an efficient way. In this dissertation we consider problems of bandwidth allocation and charging in a communication network by means of auction mechanisms.

We consider a hierarchical business model for selling bandwidth in a single link. In the top level the social planner allocates bandwidth to intermediate providers, who in turn allocate their assigned shares of bandwidth to the customers in the lower level. The social planner defines the rules of the mechanism. Our objective is the efficient overall allocation of the entire supply of bandwidth to the customers, as if the social planner were to assign this bandwidth directly. We propose an innovative mechanism comprising an appropriate auction in each level. The auctions are coordinated so that supply is exhausted at the end, as required for attaining efficiency. The bidders of both auctions have the incentives both to participate uninhibitively and to bid truthfully and efficiency be ultimately attained.

Further, we formulate business models in which the trading rules are not imposed exclusively by the social planner. In particular, we extend the hierarchical business model to a less restrictive one, in which providers are allowed to choose any payment rule to apply to their own local market, provided they conform to certain (yet fewer) restrictions. We investigate cases in which it is possible for providers to gain more profits by choosing payment rules that lead to inefficient allocations. We prove that if each customer is allowed to select his own provider on the basis of the selected payment rules, then each provider has the incentive to keep the original rule, for otherwise he would end up with no customers.

Thus far, we have assumed models for selling bandwidth in which the seller's role is restricted to defining the auction rules. We consider the problem of selling a single unit of bandwidth with the seller participating in the selling procedure too. This may be the case when the seller is also a service provider. The seller has a positive valuation for the good that is not known to the others. In order to maximize his expected profits, the seller's strategy consists of two parts: a) the choice of the auction type and b) the derivation of his own bid (which can be thought of as a reserve price that is not known to the bidders). We consider first and second-price auctions. Bidders' strategies are affected by the fact that they face one more player. Additionally, they receive a signal for the seller's valuation according to his choice of the auction type. This extra information affects their optimal strategy too. We derive both bidders' and seller's strategies in first-price and second-price auction under various information models. We compare the seller's expected profit in these two auction types and show that in some cases a first-price auction is more profitable while in other cases a second-price auction yields higher profits to the seller. We also show that this does not contradict the fact that the second-price auction with reserve price is the optimal mechanism with respect to seller profits.

We revisit the network-wide Progressive Second Price auction (PSP) proposed by Lazar and Semret. In the aforementioned mechanism, each bidder submits a consistent bid independently in each link he is interested in, taking into account the overall competition that appears in the path of his interest. We propose a new strategy for bidders taking into account the competition that appears in each link separately: we split the bidder's valuation for the demanded quantity such that each portion depicts his valuation for the same quantity in the respective link. We have carried out a wide variety of experiments that show the following: a) bidders obtain higher expected net benefit by adopting the proposed strategy than that attained with the original one. The only exception to this rule is the case where a bidder with high valuation plays first. b) the proposed strategy yields an outcome closer to the optimal social welfare than the original one.

Supervisor:
Christos Nikolaou,
Professor of Computer Science,
University of Crete

Co-Supervisor:
George D. Stamoulis,
Associate Professor of Informatics,
Athens University of Business and Economics

Εκχώρηση και Χρέωση Τηλεπικοινωνιακών Πόρων μέσω Δημοπρασιών

Μαρίνα Ε. Μπιτσάκη

Διδακτορική Διατριβή
Τμήμα Επιστήμης Υπολογιστών
Πανεπιστήμιο Κρήτης

Εκτενής Περίληψη

Η ανάπτυξη των τηλεπικοινωνιακών αγορών καθώς και η ποικιλία των προσφερόμενων υπηρεσιών έχουν καταστήσει αναγκαία τη μελέτη των μεθόδων χρέωσης στα τηλεπικοινωνιακά δίκτυα. Εκτός από την ανάγκη καταβολής τελών εκ μέρους των χρηστών, ώστε να καλύπτεται το κόστος τους, η χρέωση διαδραματίζει έναν σημαντικό ρόλο στον έλεγχο της συμφόρησης και δίνει κίνητρα για αποδοτική χρήση των δικτυακών πόρων. Οι κύριες μέθοδοι χρέωσης που έχουν μελετηθεί και εφαρμοστεί για την πώληση εύρους ζώνης είναι η επίπεδη χρέωση και η χρέωση βάσει χρήσης. Τόσο όταν εφαρμόζεται η επίπεδη χρέωση όσο και η χρέωση βάσει χρήσης, ο καθορισμός των τιμών είναι δύσκολος γιατί η ζήτηση των πελατών δεν αποτελεί γνωστή πληροφορία για τους πωλητές. Μια εναλλακτική μέθοδος χρέωσης κατά την οποία οι τιμές καθορίζονται βάσει του ανταγωνισμού των πελατών χωρίς να απαιτείται η εκ των προτέρων γνώση της ζήτησης τους, είναι η εφαρμογή δημοπρασιών. Στην παρούσα διατριβή εξετάζομε προβλήματα εκχώρησης εύρους ζώνης και χρέωσης του σε τηλεπικοινωνιακά δίκτυα μέσω δημοπρασιών.

Αρχικά, ορίζομε το πρόβλημα εκχώρησης εύρους ζώνης τηλεπικοινωνιακών δικτύων σε αγορά ιεραρχικής δομής δυο επιπέδων. Ο πωλητής του άνω επιπέδου διαθέτει διακριτές μονάδες εύρους ζώνης σε ένα σύνολο από ενδιαμέσους παρόχους τηλεπικοινωνιακών υπηρεσιών, καθένας από τους οποίους πωλεί το μερίδιο του στους δικούς του πελάτες στο κάτω επίπεδο. Θεωρούμε ότι ο πωλητής του άνω επιπέδου έχει ως στόχο την αποδοτική εκχώρηση του εύρους ζώνης στο σύνολο των τελικών πελατών και ότι αυτός καθορίζει τους κανόνες του μηχανισμού πώλησης, δηλαδή τον τρόπο της εκχώρησης και χρέωσης και στα δυο επίπεδα. Κάθε πελάτης έχει ιδιωτικά γνωστές αξιολογήσεις για τις μονάδες του εύρους ζώνης, που είναι φθίνουσες για κάθε επιπλέον μονάδα. Επίσης, για κάθε ενδιάμεσο πάροχο, η αξιολόγηση κάθε μονάδας ισούται με τα έσοδα που θα αποκομίσει από την πώληση της

συγκεκριμένης μονάδας η οποία είναι άγνωστη πριν τη συναλλαγή. Αυτή η ιδιαιτερότητα καθιστά την επίλυση του προβλήματος δύσκολη και εστιάζει στην αναζήτηση κατάλληλων δημοπρασιών για τα δυο επίπεδα κατά τέτοιο τρόπο ώστε η ανταλλαγή της πληροφορίας ανάμεσα στους παίκτες να οδηγήσει στο επιθυμητό αποτέλεσμα. Προτείνουμε ένα καινοτόμο μηχανισμό που συνίσταται από δημοπρασίες ανοικτού τύπου αυξανόμενης τιμής για κάθε επίπεδο. Αρχικά θεωρούμε ότι όλες οι δημοπρασίες είναι συγχρονισμένες μέσω ενός κοινού ρολογιού το οποίο υποδεικνύει την τρέχουσα τιμή ανά μονάδα εύρους ζώνης σε όλες τις δημοπρασίες. Οι παίκτες και των δυο επιπέδων υποβάλλουν την επιθυμητή ποσότητα σε κάθε τιμή χωρίς να γίνεται δημόσια γνωστή αυτή τους η προσφορά. Επιπλέον, κάθε παίκτης υποχρεούται να υποβάλει φθίνουσες προσφορές καθώς η τιμή αυξάνει. Ο μηχανισμός τερματίζεται όταν η ζήτηση έχει γίνει ίση με την προσφορά σε όλες τις δημοπρασίες. Κάθε παίκτης αποκτά την ποσότητα που ζήτησε στο τέλος της δημοπρασίας στην οποία έλαβε μέρος. Ο τρόπος χρέωσης διαφέρει στα δυο επίπεδα. Στη δημοπρασία του άνω επιπέδου εφαρμόζουμε τον κανόνα χρέωσης της δημοπρασίας αυξανόμενου ρολογιού με εξασφάλιση που έχει προταθεί από τον Ausubel: κάθε ενδιάμεσος πάροχος πληρώνει για κάθε μονάδα που έχει αποκτήσει, την τιμή στην οποία εξασφάλισε τη συγκεκριμένη μονάδα, δηλαδή τη χαμηλότερη τιμή κατά την οποία η μονάδα αυτή δεν ζητείται από τους αντιπάλους του. Στη δημοπρασία του κάτω επιπέδου ορίζουμε μια νέου τύπου εξασφάλιση, θέτοντας τους εξής περιορισμούς:

- Κάθε μονάδα διατίθεται από τον ενδιάμεσο πάροχο στην τιμή στην οποία την κέρδισε στο άνω επίπεδο.
- Η τιμή κάθε μονάδας την οποία πωλεί κάθε ενδιάμεσος πάροχος δεν πρέπει να είναι υψηλότερη από την τελική τιμή της δημοπρασίας του άνω επιπέδου.

Αποδεικνύουμε ότι με την εφαρμογή του παραπάνω μηχανισμού, ισχύουν οι εξής κυρίαρχες στρατηγικές: α) κάθε τελικός πελάτης αποκαλύπτει την πραγματική του ζήτηση στην δημοπρασία που συμμετέχει, και β) κάθε πάροχος αποκαλύπτει την πραγματική συνολική ζήτηση της αγοράς του στη δημοπρασία του άνω επιπέδου. Η παραπάνω στρατηγική των ενδιαμέσων παρόχων έχει ως αποτέλεσμα τον ταυτόχρονο τερματισμό ολών των δημοπρασιών που διενεργούνται και στα δυο επίπεδα. Μοναδική εξαίρεση αυτής της ιδιότητας αποτελεί η περίπτωση της μη ανταγωνιστικής αγοράς κατά την οποία ένας ενδιάμεσος πάροχος αποκτά όλες τις μονάδες εύρους ζώνης. Σε αυτή την περίπτωση, επιτρέπεται η χρέωση των μονάδων στο κάτω επίπεδο σε υψηλότερες τιμές από την τελική τιμή της δημοπρασίας του άνω επιπέδου. Ο μηχανισμός καταλήγει στη συνολικά αποδοτική εκχώρηση του εύρους ζώνης. Επιπροσθέτως, το όφελος κάθε πελάτη είναι το ίδιο με αυτό που θα του απέφερε η απ' ευθείας πώληση του εύρους ζώνης από τον πωλητή. Η ασύγχρονη διεξαγωγή των δημοπρασιών των δυο επιπέδων με τις δημοπρασίες του κάτω επιπέδου να προηγούνται, αποτελεί ένα απλούστερο τρόπο υλοποίησης του παραπάνω μηχανισμού. Περαιτέρω, διατυπώνουμε επιχειρηματικά πρότυπα που γενικεύουν το αρχικό ιεραρχικό πρότυπο. Θεωρούμε ότι οι κανόνες συναλλαγής στο κάτω επίπεδο δεν επιβάλλονται αποκλειστικά από τον πωλητή του άνω επιπέδου. Ειδικότερα, επεκτείνουμε το ιεραρχικό πρότυπο σε ένα πιο γενικό, στο οποίο οι ενδιάμεσοι πάροχοι μπορούν να επιλέξουν οποιονδήποτε κανόνα χρέωσης στην τοπική τους αγορά. Μελετάμε περιπτώσεις στις οποίες είναι δυνατό για τους παρόχους να αποκομίζουν περισσότερα κέρδη με μη αποδοτικούς κανόνες χρέωσης. Επιπλέον μελετάμε περιπτώσεις κατά τις οποίες έχει επιβληθεί στους ενδιαμέσους παρόχους ο περιορισμός της μέγιστης επιτρεπτής τιμής πώλησης σε κάθε μονάδα του κάτω επιπέδου πλην της περιπτώσεως της μη ανταγωνιστικής συμπεριφοράς. Συγκρίνουμε τον προτεινόμενο μηχανισμό με τον κανόνα ομοιόμορφης χρέωσης στο κάτω επίπεδο. Αποδεικνύουμε ότι εάν κάθε πελάτης έχει τη δυνατότητα να επιλέξει τον πάροχο του βάσει των επιλεγμένων κανόνων χρέωσης, τότε κανένας πελάτης δεν θα επιλέξει τον πάροχο που εφαρμόζει τον κανόνα ομοιόμορφης χρέωσης. Επομένως, κάθε πάροχος έχει το κίνητρο να διατηρήσει τον προτεινόμενο μηχανισμό, η επιλογή του οποίου οδηγεί σε σημείο ισορροπίας μεταξύ των διαφόρων παρόχων.

Ως εδώ, έχουμε θεωρήσει πρότυπα για την πώληση εύρους ζώνης στα οποία ο ρόλος του πωλητή περιορίζεται στον καθορισμό των κανόνων δημοπρασίας. Εξετάζουμε το πρόβλημα πώλησης μιας ενιαίας μονάδας εύρους ζώνης, με τον πωλητή να συμμετέχει στη διαδικασία πώλησης δεδομένου ότι και αυτός είναι επίσης και πάροχος. Η συμμετοχή του πωλητή είναι ισοδύναμη με την υποβολή κρυφής τιμής εκκίνησης από αυτόν. Θεωρούμε ότι κάθε πελάτης γνωρίζει τη συμμετοχή του πωλητή και έχει ιδιωτικά γνωστή αξιολόγηση για το αγαθό. Επίσης, ο πωλητής έχει ιδιωτικά γνωστή αξιολόγηση για το αγαθό. Θεωρούμε επιπλέον ότι οι αξιολογήσεις όλων των παικτών ακολουθούν την ομοιόμορφη κατανομή στο διάστημα [0,1]. Η στρατηγική του πωλητή προκειμένου να μεγιστοποιηθούν τα αναμενόμενα κέρδη του αποτελείται από δύο μέρη: α) την επιλογή του τύπου δημοπρασίας και β) την επιλογή της προσφοράς του. Εξετάζουμε δημοπρασίες πρώτης και δεύτερης τιμής θεωρώντας αρχικά ότι ο πωλητής είναι υποχρεωμένος να διενεργήσει την μια από τις δυο δημοπρασίες κάθε φορά. Οι στρατηγικές των χρηστών επηρεάζονται από το γεγονός ότι αντιμετωπίζουν έναν ακόμη αντίπαλο (τον πωλητή). Εξετάζουμε τις στρατηγικές των χρηστών και του πωλητή στις δημοπρασίες πρώτης και δεύτερης τιμής. Στη δημοπρασία πρώτης τιμής ο πωλητής αποκαλύπτει την πραγματική του αξιολόγηση για το εύρος ζώνης, ενώ στη δημοπρασία δεύτερης τιμής υποβάλει προσφορά υψηλότερη της αξιολόγησής του. Στη δημοπρασία πρώτης τιμής κάθε πελάτης υποβάλει προσφορά χαμηλότερη της αξιολόγησής του για το εύρος ζώνης, ενώ στη δημοπρασία δεύτερης τιμής αποκαλύπτει την πραγματική του αξιολόγηση. Συνεπώς, προκύπτει ότι καμία δημοπρασία δεν αποτελεί αποδοτικό μηχανισμό. Συγκρίνουμε το αναμενόμενο κέρδος του πωλητή σε αυτούς τους δύο τύπους δημοπρασίας και δείχνουμε ότι ανάλογα με την τιμή της αξιολόγησης του πωλητή σε ορισμένες περιπτώσεις η δημοπρασία πρώτης τιμής αποφέρει υψηλότερα κέρδη ενώ σε άλλες περιπτώσεις η δημοπρασία δεύτερης τιμής είναι περισσότερο κερδοφόρα για τον πωλητή. Εξηγούμε επίσης ότι αυτό δεν έρχεται σε αντίθεση με το γεγονός ότι η δημοπρασία δεύτερης τιμής με τιμή εκκίνησης είναι ο βέλτιστος μηχανισμός όσον αφορά τα κέρδη των πωλητών. Στο πρότυπο αυτό ο πωλητής αποκαλύπτει πληροφορία σχετικά με την αξιολόγηση του στους πελάτες, οι οποίοι την χρησιμοποιούν για να εξάγουν τη στρατηγική τους. Στη δημοπρασία πρώτης τιμής σύμφωνα με το πρότυπο που έχουμε ορίσει ο πωλητής δεν έχει την επιλογή του μηχανισμού, επομένως οι πελάτες εξάγουν τη στρατηγική τους βάσει της αρχικής τους γνώσης για την αξιολόγηση του πωλητή. Επομένως τα δυο πρότυπα δεν είναι στρατηγικά ισοδύναμα και αν οι πωλητές είχαν επιλογή δε θα ίσχυαν τα προαναφερθέντα αποτελέσματα για τα κέρδη τους. Είναι γνωστό από τη θεωρία ότι η δημοπρασία δεύτερης τιμής με γνωστή τιμή εκκίνησης, η οποία αποδίδει το ίδιο αναμενόμενο κέρδος στον πωλητή με τη δημοπρασία δεύτερης τιμής με κρυφή τιμή εκκίνησης αποτελεί το βέλτιστο μηχανισμό.

Τέλος, εξετάζουμε το πρόβλημα εκχώρησης εύρους ζώνης σε δίκτυο αυθαίρετης τοπολογίας με στόχο τη μεγιστοποίηση της κοινωνικής ευημερίας. Ιδιαίτερο χαρακτηριστικό της επιθυμητής λύσης αποτελεί η απαίτηση κάθε χρήστης να κερδίζει την ίδια ποσότητα εύρους ζώνης σε κάθε σύνδεσμο του μονοπατιού του. Επανεξετάζουμε την προοδευτική δημοπρασία δεύτερης τιμής (Progressive Second Price Auction) που έχει προταθεί από τους Lazar και Semret. Κατά την προοδευτική δημοπρασία δεύτερης τιμής, οι νικητές καθορίζονται βάσει των τιμών προσφορών τους και πληρώνουν για την ποσότητα εύρους ζώνης που έχουν αποκτήσει το κόστος ευκαιρίας. Με άλλα λόγια, κάθε νικητής πληρώνει το ποσό που προσφέρουν όσοι αποκλείονται από την παρουσία του. Στον προαναφερθέντα μηχανισμό, η στρατηγική που προτείνουν έχει ως εξής: κάθε χρήστης υποβάλλει την ίδια προσφορά σε κάθε σύνδεσμο του μονοπατιού του, λαμβάνοντας υπόψη το συνολικό ανταγωνισμό που εμφανίζεται σε όλους αυτούς τους συνδέσμους. Αποτέλεσμα αυτής της στρατηγικής είναι η υποβολή υψηλών τιμών από τους χρήστες μονοπατιού, τις οποίες δεν πληρώνουν (λόγω του κανόνα χρέωσης του κόστους ευκαιρίας) αλλά οδηγούν σε αρκετές περιπτώσεις σε μη αποδοτικό αποτέλεσμα. Προτείνουμε μια νέα στρατηγική με την οποία οι χρήστες λαμβάνουν υπόψη τον ανταγωνισμό που εμφανίζεται σε κάθε σύνδεσμο χωριστά, κατανέμοντας κατάλληλα τη συνολική τιμή προσφοράς για το μονοπάτι ανά τους συνδέσμους. Σύμφωνα με τη νέα στρατηγική κάθε χρήστης υποβάλλει την ίδια ποσότητα αλλά διαφορετική τιμή σε κάθε σύνδεσμο που καθορίζεται βάσει του ανταγωνισμού στο σύνδεσμο αυτό. Έχουμε πραγματοποιήσει εκτεταμένα πειράματα με τα οποία μελετάμε την αποδοτικότητα και τα καθαρά οφέλη των χρηστών υπό τις δύο στρατηγικές. Τα αποτελέσματα συνοψίζονται στα εξής: α) οι χρήστες λαμβάνουν υψηλότερο αναμενόμενο καθαρό όφελος με την υιοθέτηση της προτεινόμενης στρατηγικής από αυτό που επιτυγχάνεται με την αρχική. Μόνη εξαίρεση σε αυτόν τον κανόνα ενδεχομένως αποτελεί ο χρήστης μονοπατιού που υποβάλει πρώτος προσφορά. β) Η προτεινόμενη στρατηγική παράγει μια πιο αποδοτική εκχώρηση εύρους ζώνης από την αρχική. Η κοινωνική ευημερία υπό την προτεινόμενη στρατηγική είναι πολύ κοντά στη μέγιστη τιμή. Αντίθετα η κοινωνική ευημερία υπό την αρχική στρατηγική μπορεί να αποκλίνει σημαντικά από τη μέγιστη τιμή. Συγκεκριμένα, όσο περισσότεροι είναι οι χρήστες μονοπατιού τόσο μεγαλύτερη είναι η απόκλιση της κοινωνικής ευημερίας. Επιπλέον, όσο υψηλότερη είναι η συνάρτηση αξιολόγησης των χρηστών μονοπατιού τόσο μεγαλύτερη είναι η απόκλιση της κοινωνικής ευημερίας.

Κύριο χαρακτηριστικό όλων των προβλημάτων χρέωσης που μελετήσαμε στην παρούσα διατριβή αποτελεί η έλλειψη πληροφορίας των παικτών που εμπλέκονται στις διάφορες συναλλαγές. Η προσπάθεια κάθε συναλλασσομένου να αποσπάσει την απαιτούμενη πληροφορία από τους αντιπάλους του έτσι ώστε να μεγιστοποιήσει το ατομικό του όφελος οδηγεί πολλές φορές σε μη επιθυμητό αποτέλεσμα είτε για ένα μέρος είτε για όλο το κοινωνικό σύνολο. Η εφαρμογή της θεωρίας παιγνίων στις τηλεπικοινωνιακές αγορές διευκολύνει τη μελέτη της συμπεριφοράς των παικτών και αναμένεται να αποτελέσει ένα χρήσιμο εργαλείο για την αντιμετώπιση μελλοντικών ερευνητικών και πρακτικών προβλημάτων χρέωσης υπηρεσιών.

Επιβλέπων:
Χρήστος Νικολάου,
Καθηγητής Επιστήμης Υπολογιστών,
Πανεπιστήμιο Κρήτης

Συνεπιβλέπων:
Γεώργιος Δ. Σταμούλης,
Αναπληρωτής Καθηγητής Πληροφορικής,
Οικονομικό Πανεπιστήμιο Αθηνών

Contents

List of Figures

List of Tables

Chapter 1

Introduction

In recent years, considerable research has been conducted on the development of new pricing and trading mechanisms for network resources and services. The mandate for new charging schemes has emerged from the evolution of network technology that provides users with more complicated and customized services. Thus, appropriate charging schemes are required so that economical and engineering efficiency be reached. The high competition and the lack of information about users' demand motivate the use of auction mechanisms. The present dissertation analyzes auction mechanisms aiming to allocate bandwidth to users efficiently.

1.1 Motivation

Bandwidth is a scarce resource during congestion periods and should be utilized efficiently. To this end, it is necessary to employ economic methods to allocate resources among competing users in order to increase network performance. Prices play an important role not only because network providers need to recover their costs in order to remain in business. Different charging schemes have showed the importance of generating the right incentives to users for efficient use of network resources. Users act rationally and seek to maximize their own profits without any concern about the well-being of the society. The absence of charging rules encourages users to misuse the network resources and generate congestion, while the introduction of prices may cause a reduction in demand and control congestion. To anticipate users' behavior, the network provider should choose simple and efficient charging schemes that attract many users who are led to reveal their true valuation for the resources. This property is referred to as incentive compatibility and is essential in cases where the objective is efficient usage of resources. A complex charging scheme would make it difficult for users to realize the effect their actions have in the amount they will ultimately pay. In addition to aiming for economical efficiency, engineering efficiency is a matter of concern too. That is, it is desirable that a charging scheme is simple for the network to implement.

The simplest scheme is flat rate charging. According to this, users pay a flat fee regardless of the amount of resources they use. Extended research has been conducted for the design of usage-based charging schemes, which in contrast to flat rate charging have the following properties:

1. They reflect resource usage so that fairness among users is established. This is required for attaining economical efficiency: the resources should be shared among those users that are willing to pay the most. Under a usage-based charging scheme each user pays for what he has used, thus no one has the incentive to waste resources.

2. They deal effectively with congestion, since users restrict their demand according to their actual needs.

3. They enable the provision of differentiated services. A service is characterized by the type of traffic it handles and the quality of service required. Different services consume different amounts of resources and should be charged accordingly.

4. They give the right incentives for providers to upgrade capacity. Given that network resources are used efficiently, the network can only serve more users by increasing the amount of resources. Part of the revenues obtained by the trade can be used to expand the network capacity.

A number of researchers have proposed such schemes for best-effort services [21, 13, 6] as well as for guaranteed services [14, 10, 15]. For the latter case, charging schemes need to measure/characterize user traffic. The major objections to existing usage-based charging mechanisms are:

1. The difficulty in price determination in case of partly known demand, which applies almost always. If the network provider has no information about users' maximum willingness to

pay, he may incur losses by not setting the prices that reflect market demand. Low prices may result in decreased revenues and higher prices generate the risk to the provider of not selling the resources. In both cases, efficiency is not attained.

2. The high accounting/implementation costs. Usage-based charging schemes should be complemented by the functionality necessary to measure the amount of resources used by every provider. The design of sophisticated methods for accurate traffic measurement increases the performance of the network at the expense of an increase of implementation costs and complexity.

The balance between efficient use and simplicity is a strenuous task. Thus, almost always efficient charging approaches are complicated. Taking account of the above disadvantages, researchers have employed auction mechanisms as a means to solve the resource allocation and price discovery problems. Auction mechanisms often provide an efficient way of allocating scarce resources and achieve high value for the participants. Auction theory is suitably applied in environments where demand is not known, the information available about the value of the items is either limited or fluctuates frequently (and thus, prices need to be discovered in the process), and dynamical response to changes in market conditions is necessary. Properly designed auctions are simple, well-defined, fast and transparent mechanisms. Also, certain auction mechanisms lead users to reveal their true willingness-to-pay so that the property of incentive compatibility is met and efficiency is facilitated. Therefore, auctions are particularly suitable to the problem of network resource allocation combined with charging.

1.2 Related Work

There have been published several studies on allocating and charging network resources by means of auction mechanisms. MacKie-Mason and Varian describe in [21], a mechanism for setting the prices at network access at different priority levels; this mechanism is referred to as "smart market". In a smart market a user attaches a bid to each packet he wants to be transmitted. This bid corresponds to the user's maximum willingness to pay. All packets with bids greater than a cut-off value are admitted. Users are charged this cutoff value, which is the highest bid among packets that are not admitted. This scheme is very unstable since admission control has to be performed for each new packet. Thus, the engineering cost is very high making the smart market hard to implement. Furthermore, the uniform pricing rule that is applied does not lead users to bid truthfully.

Lazar and Semret propose in [19] the progressive second price (PSP) auction for the allocation of a divisible good (network capacity) in a single link without any a priori knowledge of demand. In particular, the PSP auction is an iterated game in which bidders are asked to submit two-dimensional bids indicating the desired quantity and the unit price. The bid price increases in time. Bidders respond to their opponents' offers until an equilibrium is reached. In the PSP auction users bid truthfully and bandwidth is allocated almost efficiently. The authors extend their approach to include bandwidth allocation in a network of arbitrary topology assuming independent PSP auctions in each link [20]. They derive an optimal strategy that combines information of all links of one's interest. A bid is then computed and submitted in each link of the path. It is claimed in [31], that equilibrium and efficiency properties still apply.

Maillé and Tuffin propose in [23] the one-shot multi-bid auction scheme for the allocation of a divisible resource in a single communication link; this scheme is closely related to the PSP auction. In particular, bidders submit a set of two-dimensional bids instead of one that is submitted in the PSP auction. This set of bids corresponds to a discretization of the bidder's utility function. The authors use an allocation and payment rule that is close to that of PSP. They prove that a bidder can not do better than reveal his true valuation and that the multi-bid auction yields an outcome that is very close to the efficient one. They extend the multi-bid auction to a special case of a network [24]. In particular, they consider a tree-like network which consists of an overprovisioned backbone network interconnected with local access networks. They adjust the multi-bid auction to allocate the same quantity of bandwidth in each link of one's path.

Courcoubetis et al. present in [9] a descending auction mechanism (MIDAS) for bandwidth allocation over paths. Bidders report their bids simultaneously and independently in each link where Dutch auctions are applied. Prices are reduced in each link at different rates and various price reduction policies are evaluated and compared experimentally. An important feature is instant allocation of bandwidth due to the property of descending prices. Experiments show that the proposed mechanism performs well in terms of efficiency. The issue of incentive compatibility is not covered completely in [9].

Reichl et al. present in [27] two extensions of the Generalized Vickrey Auction that allow for a dynamic pricing of flow-based Internet traffic covering more than one Internet Service Provider. In particular, the authors of [27] deal with Delta Auctions which provide a solution for the delay problem caused by multiple auctions necessary for setting up a connection crossing multiple Internet Service Providers. Moreover, the authors of [27] present a multi-provider auction model (CHiPS) based on the idea that the holders of already running connections are preferred than newly arriving bidders, so that their connections are not interrupted in the immediate future.

Reichl et al. deal with auctions for multi-period sessions and multi-link connections in the Internet in [26]. The authors of [26] propose the Second-chance Auction Mechanism (SAM). This is a combination of the MIDAS auction of [9] with the CHiPS mechanism of [27] in which users who have received allocations in a certain auction (in time) should be allocated bandwidth in the subsequent auction as well. The authors of [26] claim that this results in improved efficiency even for the network case.

Lalis et al. propose in [18] a trading mechanism based on a continuous double auction. This mechanism is employed for matching application requests to the services offered over an ATM switch. The authors of [18] investigate the performance and robustness of the mechanism through a series of experiments. In particular, end-to-end quality of service is achieved when clients compete in multiple autonomous resource markets, for a fixed period of resource usage.

1.3 Problems Analyzed and Our Contribution

In this section we summarize the problems analyzed in this dissertation as well as our contribution.

Bandwidth allocation through two hierarchical levels. We formulate a new problem of allocating bandwidth through intermediaries, in a single link. In the top level the social planner allocates bandwidth to intermediate providers, who in turn allocate their assigned shares of bandwidth to the customers in the lower level. The social planner defines the rules of the mechanism. Our objective is the efficient overall allocation of the entire supply of bandwidth to the customers, as if the social planner were to assign this bandwidth directly. We propose an innovative mechanism comprising an appropriate auction in each level. The auctions are coordinated so that supply is exhausted at the end, as required for attaining efficiency. The bidders of both auctions have the incentives both to participate uninhibitively and to bid truthfully and efficiency be ultimately attained. Our approach has the important property that users pay the same despite the existence of intermediaries.

Furthermore, we formulate business models in which the trading rules are not imposed exclusively by the social planner. In particular, we extend the hierarchical business model to a less restrictive one, in which providers are allowed to choose any payment rule to apply to their own local market, provided they conform to certain (yet fewer) restrictions. We investigate cases in which it is possible for providers to gain more profits by choosing payment rules that lead to inefficient allocations. We prove that if each customer is allowed to select his own provider on the basis of the selected payment rules, then each provider has the incentive to keep the original rule, for otherwise he would end up with no customers.

Seller participation in a single-unit auction. It is very common in telecommunication markets for the owner of bandwidth, to have a positive valuation by keeping part of it for himself. For example, the owner may be an ISP too, as his customers are. We address the problem of selling a single unit of bandwidth (i.e. the entire capacity of a link) with the seller participating in the selling procedure too. The seller has a positive valuation for the good that is not known to the others. In order to maximize his expected profits, the seller's strategy consists of two parts: a) the choice of the auction type and b) the derivation of his own bid, which can be thought of as a reserve price that is not known to the bidders. We consider first and second-price auctions. Bidders' strategies are affected by the fact that they face one more player. Additionally, they receive a signal for the seller's valuation according to his choice of the auction type. This extra information affects their optimal strategy too. We derive both bidders' and seller's strategies in first-price and second-price auction under various information models. We compare the seller's expected profit in these two auction types and show that in some cases a first-price auction is more profitable while in other cases a second-price auction yields higher profits to the seller. We also show that this does not contradict the well-known fact that the second-price auction with reserve price is the optimal mechanism with respect to seller profits.

Bandwidth allocation in a communication network. We revisit the network-wide Progressive Second Price auction (PSP) proposed by Lazar and Semret in [20]. We have examined thoroughly the efficiency of the mechanism and have discovered a major deviation from the optimal social welfare. We argue that this is caused due to the strategy proposed for path bidders. According to this strategy, the optimal quantity to be submitted is the largest one such that the bidder's marginal valuation is just greater than the total market price. The total market price is the sum of the prices at the different links. The optimal price to be submitted in each link of the bidder's path, is then his marginal valuation at the optimal quantity, which reflects the total demand of the market. We propose a new strategy for path bidders taking into account the competition that appears in each link separately: the bid price is now defined to be the minimum price at which the bidder obtains the demanded quantity in a specific link. Thus, the bid prices differ in the various links, reflecting the demand of each link as opposed to the total demand, which is the case with the strategy of [20]. We have carried out several experiments that show the following:

- bidders obtain higher expected net benefit by adopting the proposed strategy than that attained with the original one. The only exception to this rule is the case where a bidder with high valuation plays first and essentially drives the auction to immediately terminate.
- the outcome of the proposed strategy is always very closer to the optimal social welfare, than the original one.

1.4 Outline

The remainder of the dissertation is organized as follows. In Chapter 2 we introduce basic concepts of auction theory and review the most important auction mechanisms for a single good, for many identical goods, and for many heterogeneous goods. Basic results and comparisons of these mechanisms are also discussed at the end of each section.

In chapter 3 we consider the problem of allocating bandwidth in a single link, through intermediaries. First, we define the problem and discuss about the requirements a solution must satisfy in order to have our objective fulfilled. We highlight basic issues of the problem by analyzing the case of allocating a single unit of bandwidth. We then describe the hierarchical auction mechanism to allocate bandwidth in two levels, and we prove the basic properties of the mechanism that include bidders' strategies and efficiency, and discuss about alternative implementations of the mechanism and other practical issues. Finally, we introduce various market-based models in which intermediaries (respectively customers) have been given more flexibility for their charging rules (respectively which intermediary to associate with). We discuss about the applicability of our mechanism and the related consequences.

In chapter 4 we consider the problem of seller participation when a single good is being auctioned. First, we review the problem of auctioning a single good by imposing a reserve price. Furthermore, we derive bidders' and sellers' strategies in a first-price auction and in a second-price auction with the seller participating as if he were one of the bidders. We then compare the first-price auction and the second-price auction in terms of seller expected profit and discuss about a paradox that is concerned with the optimal auction for selling one single good.

In Chapter 5 we deal with the problem of allocating bandwidth in a network of arbitrary topology. First, we revisit the network-wide PSP auction and discuss the weaknesses we have investigated. We then define a new strategy for the PSP auction. We describe the set of experiments we have carried out to comparatively assess the two strategies, and present the corresponding results.

Finally, in Chapter 6 we provide some concluding remarks and discuss directions for future work.

Chapter 2

Overview of Auctions

2.1 Introduction

An auction is a method for selling a set $\mathcal{C}$ of goods to a set $\mathcal{N} = \{1, \ldots, N\}$ of buyers in which the *auctioneer* (seller or a third party) sets the rules and the *bidders* (potential buyers) place their offers. The outcome of the auction is specified by the allocation and the payment rules, which determine the winners and their payments.

The seller is characterized by his *cost function* and his *utility function* $u : \mathcal{C} \rightarrow \Re$, where $u(q)$ denotes how much quantity q is worth to him if it is not sold. In the present dissertation, we consider only cases where the set of goods is fixed and commonly known to all players, and the seller's cost function equals zero. Each bidder i is characterized by his utility function (or valuation) $\theta_i : \mathcal{C} \rightarrow \Re$, where $\theta_i(q)$ denotes the maximum willingness to pay of bidder i for buying the whole quantity q. In case where q takes a continuum of values, the *derivative* θ_i' of the utility function is the marginal utility. $\theta_i'(q)$ denotes the price at which bidder's i demand equals q; that is the inverse demand function. Recall that the demand function $q(p)$ at price p for a bidder, denotes the quantity the bidder is willing to buy if the price is p. On the other hand, if the set $\mathcal{C}$ of goods is discrete, the marginal valuation $\theta_{i,j}$ of bidder i for the j^{th} additional unit he may acquire, is the maximum amount he is willing to pay for that unit. In this case, the valuation of bidder i for acquiring q units equals the sum of the q marginal valuations, that is $\theta_i(q) = \sum_{j=1}^{q} \theta_{i,j}$. The marginal valuations form the respective inverse demand function. The following example clarifies these issues in the case where identical indivisible units of a single good are being auctioned.

Example 2.1.1

Consider a set of four identical units of a good. Let $\vec{\theta}_i' = (\theta_{i,1}, \ldots, \theta_{i,4}) = (8, 5, 4, 2)$ be the vector of bidder's i marginal valuations. Then, for the first unit he is willing to pay 8 units of money, for the second unit (provided that he has won the first one) he is willing to pay 5 units and so on. The total valuation of acquiring c ($c \leq 4$) units is given by the sum $\sum_{j=1}^{c} \theta_{i,j}$ of the first c marginal valuations. Following the notation above, the vector $\vec{\theta}_i'$ forms the inverse demand function. The inverse demand function and the utility function of bidder i are depicted in figure 2.1. ▲

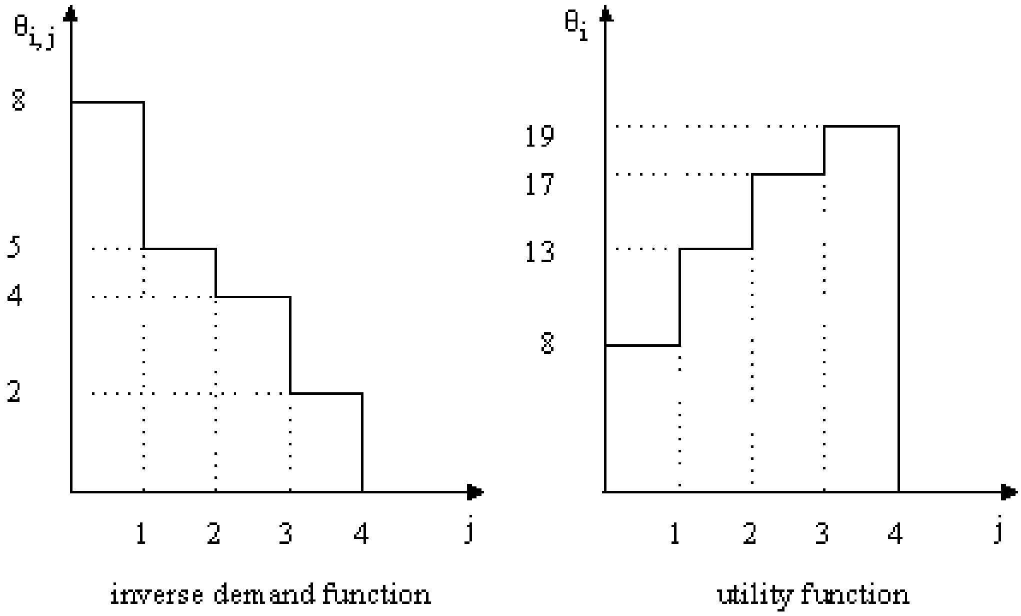

Figure 2.1: Marginal valuations and utility function of bidder i in Example 2.1.1

Another feature that characterizes a player (seller or bidder), is the level of *risk aversion*. We introduce the Von Neumann and Morgenstern utility function to explain these notions; for a thorough study see [3]. Let $V : \Re \rightarrow \Re$ denote the Von Neumann and Morgenstern utility function of a player. V expresses the player's preferences over the expected profits he gains from the auction. If V is a concave function of the expected profits, this means that each additional unit of money is worth less to the player than its previous one. Thus, the player would prefer an amount equal to the expected utility $E(V)$ of all possible profits than participating in the auction, if he had the opportunity to choose. This player is called risk averse. A risk-averse player seeks to maximize the utility of his expected profits. If V is a linear function of the expected profits, then the player is indifferent between participating or being given an amount equal to the expected utility. This player is called risk neutral and seeks to maximize his expected profits.

In order to analyze and design an auction, one should take into consideration the rules, the players' utility functions and the *information* available to each player about his rivals' characteristics. The two extreme cases, considering bidders' valuations, usually used in auction theory, are:

- the *private values* model in which each bidder's valuation is privately known to him and not affected by the others' valuations and
- the *common values* model in which all bidders have the same unknown valuation which is only observed, after the end of the selling procedure.

In any case, bidders' valuations are not known to the seller and to their rivals, thus, an auction is considered a game of incomplete information. In fact, the uncertainty about bidder's i valuation can be described by a probability distribution: the seller and the rivals of bidder i do not know bidder's i valuation but know (common knowledge) the distribution this valuation is drawn from. When these distributions are the same for all bidders, then bidders are said to be *symmetric*.

2.2 Analysis of an Auction: Objectives and Strategies

Though auctions are easy to organize and implement, it is difficult to design and analyze an auction such that the desired *incentives* for the bidders are induced. Thus, the auctioneer should specify the auction so that his objectives are achieved, taking into consideration how bidders will play in such an auction. Given the rules and the information available, each bidder selects his own actions, that is his *strategy*, such that his own objective is attained.

Among the objectives a seller might pose, is *profits maximization* and *social welfare maximization* (*efficiency*). In the first case, the seller wants to select the auction that maximizes his profit subject to a quantity constraint. The profit for a seller equals the revenues he obtains from the sold quantity plus the utility he receives from the unsold quantity. Since both revenues and quantity finally sold are not known in advance but are assumed to depend on bidders' distributions of their valuations, the seller seeks to maximize the expected profits. Auctions having this property are referred to as *optimal auctions*. In the case of social welfare maximization, the seller is a non-profit organization or a government and may also be referred to as social planner. He wants to allocate (more suitable verb than "sell") the quantity fairly to those bidders who value it most. That is, to maximize the social welfare SW, which measures the well-being of the society consisting of all parts that are involved in the specific trade. Social welfare is given by the sum of profits of all players (seller and bidders). The profit for a bidder is called *net benefit*. For a won quantity q, it is given by $NB(q) = \theta(q) - c(q)$, where $c(q)$ is the amount the bidder is charged for obtaining quantity q. As a result, the social welfare associated with the vector $\vec{q} = (q_1, \ldots, q_N)$ of allocated quantities for the N bidders is given by:

$$SW(\vec{q}) = \sum_{i \in \mathcal{N}} NB_i(q_i) + u(q_s) + r(q_b) = \sum_{i \in \mathcal{N}} \theta_i(q_i) + u(q_s), \tag{2.1}$$

where q_s is the unallocated quantity, $q_b = q_1 + \cdots + q_N$ the allocated quantity and $r(q_b)$ the revenues the seller obtains from selling quantity q_b. Note that SW is independent of r. Thus, if all goods are finally sold, the social welfare equals the sum of the bidders' valuations for the won quantity. Note that in the case of the social welfare maximization objective, we do not take the expected value as in the case of profit maximization. Even though θ_i's are not publicly known, the objective is to find the allocation of the goods that produces the efficient outcome among all feasible outcomes. In subsequent analysis, we will see that the objectives of profits maximization and social welfare maximization, are generally in conflict.

As soon as the auction's rules are set by the seller, the bidders are ready to calculate their own strategies. A bidder's strategy is defined as the bidding actions he takes at each time he has to make a decision during the whole procedure and for each state of information the bidder may have at such times. Auction theory, as part of game theory, is based on the assumption that each bidder acts rationally. That is, he selects a strategy such that his utility of the expected net benefit [1] be ultimately maximized, given the auction's rules. Information exchange among the different parties of an auction game plays a significant role in analyzing strategic behaviors. The information (belief) each bidder has about others' utility functions is part of the assumptions

[1] In the case of a risk neutral bidder, rational behavior enforces maximization of the expected net benefit.

about the specific situation in which the auction is applied, whereas the information bidders have or may acquire about their opponents' bidding is part of the auction's rules. The following basic strategies will be considered in the next sections:

- *Truthful bidding* is the strategy in which the bidder reveals his true utility function through his offers.
- *Bid shading* is the strategy in which the bidder declares a lower price than his valuation for a quantity; alternatively, we have *demand reduction*, where the bidder declares less quantity than his real demand at a given price.

Analyzing one's rational behavior is far more complicated than considering only the auction's rules and the information state. A bidder's strategy influences and is influenced by the strategies chosen by the others. Thus, in order to determine the best strategy, a bidder has to take into consideration what the others' strategies will be and that the others will make assumptions for his own strategy too and so forth. A set of strategies, one for each bidder, is called a *Nash equilibrium* if no single bidder has the incentive to unilaterally deviate from his strategy, if he knows that all the other bidders do employ these strategies. An even stronger property of a strategy is that of dominance. A *dominant strategy* is an optimal strategy for a bidder irrespectively of the others' strategies. The special case where truthful bidding constitutes a Nash equilibrium is called the *incentive compatibility* property. Throughout this dissertation, bidders will be assumed to select their strategies without making collusions. Thus, auctions will be analyzed as non-cooperative games.

2.3 Classification of Auctions

There are various criteria to classify auctions. We distinguish among single-unit, multi-unit and multi-object auctions, depending on the number and nature of the goods to be sold. In multi-unit auctions, identical (homogeneous) units are being auctioned such as for example units of bandwidth in a communication link. On the other hand, multi-object auctions refer to heterogeneous goods that might be complements or substitutes. A bundle of goods is said to be complementary to a bidder, if his valuation for the bundle is higher than the sum of valuations of each good of the bundle taken separately. For example, consider the problem of selling bandwidth in a path of a communication network. A bidder wishes to obtain bandwidth in every link comprising his desired path, otherwise it is useless to him. Similarly, a bundle of goods is said to be substitute to a bidder, if his valuation for the bundle is less than or equals the sum of valuations of each good of the bundle taken separately, that is, goods of the bundle may be substituted by others in the same bundle. Identical units of a good are perfect substitutes.

We further distinguish between oral and written auction formats. An oral auction is a progressive procedure in which bidders have the chance of making counter-offers and bidding information may be released, while a written auction is a one shot bidding process in which bidders submit a sealed bid, and have no chance of responding back.

There have been carried out thorough studies about auction design and very important results have been established, especially in the single-unit case. In the sequel, we overview some well-known auctions for selling a single good or multiple goods, some of which can be successfully used to sell bandwidth in a communication link either seen as a set of identical and indivisible units, or a single arbitrarily divisible unit.

2.4 Single-Unit Auctions

There are many different auction types used to sell a single good, each one of which is best suited in different circumstances. An auction is said to be *standard* if the good is sold to the bidder who submitted the highest bid. We briefly describe the most commonly used standard auction types and discuss about their properties. A more detailed exposition can be found in [35].

English Auction. The English auction is an open auction performed in consecutive rounds. The auction starts at a reserve price, which is the lowest acceptable price set by the auctioneer. The bidders raise the price in each round until the standing highest bid is no more improved. The bidder willing to pay the most is the winner and pays his bid, unless nobody submits a bid higher than the reserve price. The best strategy for each bidder is to increase his bid in each round by a small amount more than the previous standing bid up to his valuation and then stop. This is optimal, since bidding a higher amount would not increase his probability of winning but would increase the payment, had the auction terminated at that time. It is important to mention that the final price is determined by the second highest valuation of all bidders: the last bid is submitted by the bidder with the highest valuation and equals the second highest valuation plus a small increment. This type of auction is highly susceptible to collusion rings, gives great influence to the auctioneer and is the most emotional and competitive of auctions. Bidders' presence is required and the seller does not necessarily receive the maximum value (not optimal). A principal property of an ascending auction is that it reveals information to the bidders through the process of bidding.

Dutch Auction. The Dutch auction is an open auction too. The auctioneer proposes an initial very high price, which then descends progressively. The first bidder who calls out that he wants the unit, is the winner and pays the current price. Regarding the strategy, each bidder has to decide at which price he intends to claim the good before entering the game, since no information is released during the auction. The only information revealed is the final price and the winner, which are given at the end, when they are useless. A high bid (up to one's valuation) increases the probability of winning, whereas a lower bid induces higher net benefit in case of winning. The optimal bid shading for each bidder is based on his valuation for the good and his beliefs about the others' valuations. The Dutch auction is less susceptible to collusion rings and the auctioneer has almost no influence.

First-price sealed-bid Auction. Each bidder submits a sealed bid without knowing others' bids and the unit is awarded to the highest bidder. The winner pays his bid. This type is strategically equivalent to the Dutch auction; that is each bidder performs bid shading according to his beliefs about his rivals' valuations. The following example presented in [16], employs a first-price auction and shows how the optimal bid shading can be determined.

Example 2.4.1

Consider a set of N risk-neutral and symmetric bidders competing for a good by means of a first-price sealed-bid auction. Each bidder i has a valuation X_i for the good. Each X_i is independently drawn from the uniform distribution on the support $[0, 1]$. Bidder i knows the realization (i.e. is the value) x_i of X_i and that his rivals' valuations are uniformly distributed. Then the symmetric equilibrium strategy for each bidder i is given by:

$$\beta(x) = \frac{N-1}{N} \cdot x. \tag{2.2}$$

A sketch of the proof is as follows: suppose that each bidder but bidder i applies strategy β and bidder i has valuation x and bids b. The winning probability of bidder i equals $P[\max_{j \neq i} \beta(X_i) < b] = P[\max_{j \neq i} X_i < \beta^{-1}(b)] = (\beta^{-1}(b))^{N-1}$. Thus, bidder's i expected net benefit equals $NB_i = (\beta^{-1}(b))^{N-1} \cdot (x - b)$. We take the derivative with respect to b, and solve the differential equation $(NB_i)' = 0$. We obtain $b = \beta(x) = \frac{N-1}{N} \cdot x$.

We note that, even though bidders lie about their true valuation, the outcome is efficient in the symmetric case, since the same bid shading is performed by each bidder. Thus, the order of their bids coincides with the order of their valuations. ▲

Vickrey Auction. Each bidder independently submits a sealed bid without knowing others' bids and the unit is awarded to the highest bidder. The winner pays the second-highest bidder's bid. It is a dominant strategy for each bidder to reveal his true valuation: a bidder who bids more than his valuation faces the risk of winning the good at a price higher than his valuation, while bidding less increases the probability of not winning. In either case, the lowest possible price reduces to the second highest valuation, which is the same if truthful bidding is applied. Therefore, the Vickrey auction is an incentive compatible mechanism.

Principal Properties of Standard Single-Unit Auctions. We summarize some interesting results that arise by comparing the above four auction types with respect to seller's revenues, a complete description and justification of which can be found in [35] and [16]. The Dutch auction is strategically equivalent to the first-price sealed bid auction, since both types release no information about bidders' valuations before the end. Under private values, the English and the Vickrey auctions result in the same outcome (in both types the second highest bidder determines the winner and his payment), but are not strategically equivalent: in an English auction, bidders have the opportunity to respond to their rivals' bids. These results were established by Vickrey in [33], who additionally showed that all the above standard auction types are *revenue equivalent*, given that: bidders have private values, independently distributed, are symmetric, risk neutral and the seller is risk neutral too (SIPV model). Myerson in [25] extended the above revenue equivalence result to show that *all* auctions that award the item to the bidder with the highest valuation and lead to the same bidder participation, are revenue equivalent. However, if the seller is risk averse or the bidders are risk averse, then the Dutch auction yields higher revenues to the seller than the English auction. A risk averse bidder prefers a higher probability of winning at a cost of a lower net benefit. Hence in a first-price auction his bid shading is expected to be less aggressive, thus yielding higher revenues. If the assumption of symmetric bidders is relaxed, first-price auctions and second-price auctions are no more revenue equivalent: bid shading in a first-price auction might not be the same for all players, resulting in a different (probably not efficient) outcome than that of the second-price auction, in which truthful bidding still applies. In [16], appropriate examples are presented that show that no ranking, with respect to revenues, can be obtained for the two auction types, when asymmetric bidders take part. Under the common values model, the English auction is no more equivalent (in terms of outcome) with the Vickrey auction: information released in an English auction helps bidders make better estimates about the good's valuation.

Optimal Auctions. Even though efficiency is attained in any single-unit auction under the SIPV model (i.e., the good is sold to the bidder who values it most), these auctions are *not* optimal. Efficiency and revenue maximization are indeed in conflict. The best the seller could do, in terms of revenues, is to receive the maximum valuation among all bidders' valuations, were he aware of bidders' valuations for the good. When applying auctions, he can at least receive the second highest valuation as mentioned above, still this is not optimal. The optimal auction among all single-unit auctions in the symmetric case is established by Myerson in [25]. This is a second-price auction, modified by introducing a reserve price that is higher than the seller's valuation and depends on the distribution of bidders' valuation. The winner is the bidder that has submitted the highest bid provided that it is higher than the reserve price, and pays the maximum of the second highest bid and the reserve price. Imposing a reserve price, the seller may lead to the reduction of the number of participants and face the risk of keeping the good even though there is a bidder willing to pay more than the seller's valuation (the optimal auction may not be efficient). This effect is outperformed by the probability of the reserve price being higher than the second highest valuation, thus resulting in a higher expected value of the revenue. Optimal auctions are further studied in [28] and [4].

2.5 Multi-Unit Auctions

In this section, we briefly discuss about multi-unit auctions and summarize some basic properties. A more detailed discussion can be found in [8] and [16]. Multi-unit auctions of discrete goods can be either *sequential*, in which each good is sold one after the other or alternatively *simultaneous* in which all goods are auctioned at the same time. Furthermore, we characterize multi-unit auctions as *discriminatory* (pay-your-bid) or *uniform* according to the payment rule that is used. Examples for both types are given below. We conclude the discussion of multi-unit auctions by describing an auction for an infinitely divisible good such as bandwidth. All auctions that will be analyzed in this section are *standard auctions*; that is, all available units are awarded to the bidders with the highest bids.

2.5.1 Sealed-Bid Auctions

We summarize three types of sealed-bid auctions to sell C identical units of a good, namely the discriminatory, the uniform and the generalized Vickrey auction as presented in [16]. All these auctions apply the same allocation rule but differ in the payment rule, and thus, in the bidding behavior. We consider the private values model and that bidders' marginal valuations are *non-increasing*: the bidder assigns a lower value to each extra unit he obtains.

Uniform Auctions. Each bidder submits C bids indicating the valuation he assigns to each additional unit (i.e., his marginal valuations). These bids form his demand function that denotes how many units he is willing to buy at each price. Winner determination is performed by adding all bidders' submitted demand functions so that the aggregate demand function is formed. The aggregate demand function denotes how many units, bidders are willing to buy in total at each price. The price at which the aggregate demand function intersects with supply is the *market clearing price*. All bids above that price are winning bids. In other words, all bidders' submitted bids are ordered and the C highest ones are the winning bids. The market clearing price is any price lying between the highest losing bid and the lowest winning bid. Each bidder pays the market clearing price for each unit he obtains (uniform pricing). Figure 2.2 illustrates the aggregate demand function. The lowest winning bid equals p_2 and the highest losing bid equals p_1. The market clearing price may take any price between p_1 and p_2.

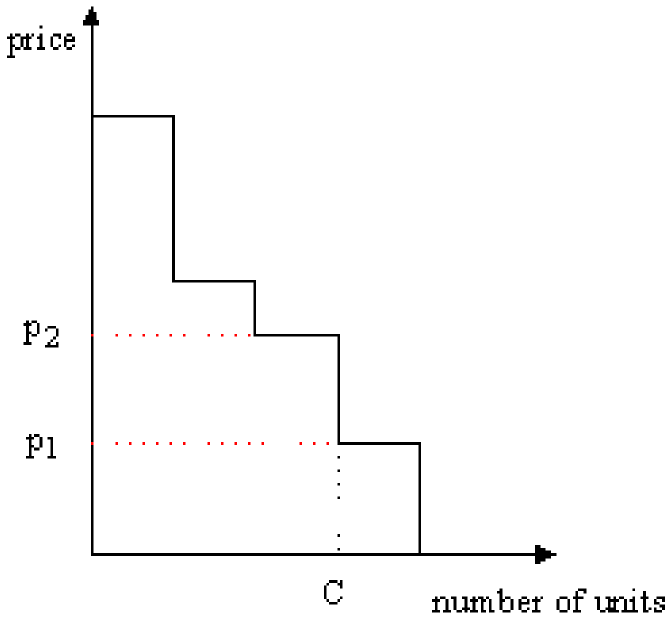

Figure 2.2: Aggregate demand function

The bidding behavior in a uniform auction is rather complicated. With uniform pricing bid shading is enforced: each bidder has the incentive to submit a lower bid than his marginal valuation for each additional unit. This way, he manages to reduce the market clearing price (demand now equals supply at a lower price), thus, paying less for each unit won. Even though he thus obtains less units, there is a market clearing price at which he incurs a higher net benefit. Equilibrium strategies (i.e. the optimal degree of bid shading) that induce such prices do exist but are difficult to compute. Nevertheless, interesting properties regarding these equilibrium strategies have been proved (see [16]). First, bids cannot be higher than bidders' respective marginal valuations. Second, it is a weakly dominant strategy for any bidder to submit his true marginal valuation for the first unit. Bid shading is then performed in the subsequent units. In particular, bid shading becomes more aggressive in the last units. This differential bid shading causes uniform auctions to yield inefficient outcomes. In the restricted model where each bidder wishes at most one unit, truthful bidding is a weakly dominant strategy and the auction is efficient. In the case of auctioning one unit of a good, a uniform auction reduces to the Vickrey auction provided that the market clearing price is taken to be the highest losing bid (i.e., the second highest bid).

Discriminatory Auctions. In a discriminatory auction winners are determined similarly with the uniform auction but a different payment rule is applied. Each bidder pays his bid for each unit he wins. Bid shading is the only way for a bidder to extract a positive net benefit. Again the trade off between increasing the probability of winning or increasing the expected net benefit is considered. This auction is inefficient too. In the case of one unit, a discriminatory auction reduces to the first-price auction discussed in the previous section.

Generalized Vickrey Auction (GVA). Vickrey proposed in [33] an efficient mechanism to allocate C identical units to players with non-increasing marginal valuations. Each bidder submits C bids in decreasing order that represent his willingness to pay for each additional unit he might win. The C highest bids are the winning bids. A bidder that has won C_j units, is charged the total amount of the C_j highest rejected bids that have been placed by the other bidders (and not including his own). This means that each winner pays the *social opportunity cost*, which equals the extra value that would have been generated to the other bidders in his absence. Truthful bidding (i.e. revealing the truthful demand function) is a weakly dominant strategy yielding an efficient outcome. In the case of one unit, GVA reduces to the Vickrey auction. The following example illustrates the uniform and Vickrey payment rules.

Example 2.5.1

Consider two bidders competing for two identical units, with privately known and independent marginal valuations shown in Table 2.1. If we apply the GVA, bidders are truthful and the efficient

Bidder 1	Bidder 2
$\boldsymbol{\theta_{1,1} = 8}$	$\theta_{2,1} = 4$
$\boldsymbol{\theta_{1,2} = 5}$	$\theta_{2,2} = 2$

Table 2.1: Marginal valuations in Example 2.5.1

outcome is attained: bidder 1 wins both units. He pays the two rejected bids 4+2=6. In the case of applying the uniform auction, each bidder bids truthfully for his first respective unit, while shades his bid to zero for the second one. These are equilibrium strategies that yield an inefficient outcome: each bidder wins one unit. The market clearing price equals zero, so both bidders pay zero. ▲

2.5.2 Open Auctions for Selling Identical Units

Ascending Auctions. Ascending auctions are very well explained in [12]. We provide a brief description herein. In the general form of an ascending auction mechanism for the case of identical units, winners and prices are determined in an open competition among bidders. All units are being auctioned simultaneously and bidders willing to pay the most are the winners. Specifically, each bidder submits a bid of the form $s = (q, p)$ in each round, where q is the number of units he would be willing to buy at price p. This bid replaces his previous one. Each bidder is assumed to have a non-increasing demand curve. Thus, the activity rule that is imposed for each bidder is a) to increase his bid price and b) decrease the demanded quantity in each round. The set of all bids forms the aggregate demand function. In the initial round, aggregate demand is high (higher than supply). In subsequent rounds demand falls as price increases until demand equal supply, where the auction terminates. The price at which demand equals supply is the market clearing price. After the end, the bids are ordered from high to low bid prices and the bidders who have bidded above or at the market clearing price are the winners. Each winner obtains the units he demanded at his final bid except for the last bidder who obtains the remaining units (which may be less than the quantity he has demanded). Two alternative payment rules may be applied: uniform pricing under which winners pay the market clearing price for each unit they win, or pay-your-bid pricing (a case of discriminatory pricing) in which each winner pays a different price, namely his bid.

An alternative simpler version of the ascending auction described in [12], is the Ascending clock auction. A clock that indicates the current price increases in each round. At each price, each bidder submits the quantity he is willing to buy, instead of a quantity and a price as bidded in the general format. A bidder cannot increase his quantity as the price increases. Bidding continues until the total quantity requested by all bidders equals supply. Winners and payments are determined as above.

The bidding behavior in ascending auctions with uniform pricing is similar as in the sealed-bid uniform auction in the private values model. Instead of the term bid shading, demand reduction is equivalently used: each bidder has the incentive to misreport his demand function and bid for a quantity less than his true one at each price. Thus, the two auction formats are equivalent with respect to their outcomes. In the case of common values, though, ascending auctions release information throughout the process and equivalence is no more valid. With pay-your-bid pricing, bid shading is inevitable too. Both pricing rules are studied and compared with respect to seller's revenues in [2]. The ranking is ambiguous depending on the bidders' demand curves. Both approaches discussed above lead to inefficient outcomes.

Note that the probability of a simultaneous decrease in demand by two or more bidders at any price is negligible. This is due to the fact that bidders are assumed to have marginal valuations taken from a continuous set, thus the probability that two or more bidders have the same marginal valuation for a unit, equals zero. However, it is possible for a certain bidder to reduce demand by two or more units of the good, since each bidder is not assumed to have strictly decreasing marginal valuations. Thus, it is possible that the final demand be less than supply in an ascending auction. For simplicity reasons, we henceforth assume that demand does equal supply at the final price of any ascending auction that is considered in this dissertation. In the case where demand would be less than supply, the remaining unallocated units of bandwidth would be assigned to providers according to a rationing rule; for example the rule used in [1].

Ascending Clock Auction with Clinching (ACC). An ascending auction referred to as Ascending Clock Auction with Clinching was proposed by Ausubel in [1]. This can be applied to sell C units of a good (e.g. bandwidth) efficiently. The auctioneer starts at a reserve price that continuously increases as the auction proceeds by means of a price-clock. At each price, the bidders simultaneously submit the desired quantity which is required to be a non-increasing function of price. Thus, aggregate demand (sum of bids in each price) decreases and the auction terminates when it equals supply. Each bidder wins the quantity he demanded at the final price but is not charged this unit price. Instead of applying a uniform-pricing rule that would result in demand reduction, each bidder is charged for each unit he wins the standing price at which he *clinched* (i.e. he was to be allocated with certainty) this same unit. That is, a bidder wins a unit at the first time this unit is *not* claimed by his rivals; i.e., whenever there is a decrease in his rivals' aggregate demand, provided that it is less than supply. Under this payment rule, assuming pure private values, non-increasing marginal utilities and no information on others' bids (zero-information model), truthful bidding by every bidder is a weakly-dominant strategy that yields an efficient outcome. Alternatively, each bidder may have full information on others' bids at each price (full bid information) or may know the sum of bids of his rivals at each price (aggregate bid information). In the case of full bid information or aggregate bid information, truthful bidding by every bidder constitutes a Nash equilibrium. GVA is outcome equivalent to the ACC auction with no bid information: they have the same pricing rule and bidders follow the same strategy. But the ACC auction does have an advantage over the GVA auction. The ACC auction reveals only a portion of the winners' utility functions; their maximum willingness to pay is still preserved private. Example 2.5.2 shows how ACC is applied.

Example 2.5.2

Suppose there is an amount of $C = 8$ units of bandwidth that is allocated to five bidders, with marginal valuations [2] shown in Table 2.2.

Bidder 1	Bidder 2	Bidder 3	Bidder 4	Bidder 5
$\boldsymbol{\theta_{1,1} = 10}$	$\boldsymbol{\theta_{2,1} = 17}$	$\boldsymbol{\theta_{3,1} = 15}$	$\boldsymbol{\theta_{4,1} = 12}$	$\boldsymbol{\theta_{5,1} = 19}$
$\theta_{1,2} = 4$	$\boldsymbol{\theta_{2,2} = 11}$	$\theta_{3,2} = 7$	$\theta_{4,2} = 8$	$\boldsymbol{\theta_{5,2} = 18}$
$\theta_{1,3} = 2$	$\theta_{2,3} = 1$	$\theta_{3,3} = 5$	$\theta_{4,3} = 3$	$\boldsymbol{\theta_{5,3} = 16}$

Table 2.2: Marginal valuations in Example 2.5.2

We apply ACC with the price starting at price 1. Tables 2.3 and 2.4 show the bidding and the clinching process respectively. Note that although the price is continuous, we only present the price values at which new bids arise. The auction terminates at price 8 with bidder 1 winning 1 unit, bidder 2 winning 2 units, bidder 3 winning 1 unit, bidder 4 winning 1 unit and bidder 5 winning 3 units, which is the efficient allocation. Bidder 1 clinches his first unit at price 8, since his opponents' demand at price 8 is 7 units, whereas the supply is 8. Thus, bidder 1 obtains the remaining unit (8-7) with certainty at price 8, which is his payment for that unit. Similarly, bidder 2 clinches his first and second unit at prices 7 and 8 respectively. Bidder 3 clinches his first unit at price 8, bidder 4 clinches his first unit at price 7 and bidder 5 clinches his first, second

[2] Recall that $\theta_{i,j}$ denotes the marginal valuation of bidder i for the additional unit j.

	Bidding Process				
Price	Bidder 1	Bidder 2	Bidder 3	Bidder 4	Bidder 5
1	3	2	3	3	3
2	2	2	3	3	3
3	2	2	3	2	3
4	1	2	3	2	3
5	1	2	2	2	3
7	1	2	1	2	3
8	1	2	1	1	3

Table 2.3: Ascending Clock Auction with Clinching in Example 2.5.2: bidding process

and third unit at prices 5, 7 and 8 respectively. Note that even though total demand has dropped by one unit from price 5 to price 7, three bidders have clinched one unit each. The two last rows of Table 2.4 show the payments and the net benefits. Bidder 1 pays a total amount of 8 units of money, bidder 2 pays a total amount of $7 + 8 = 15$ units of money, bidder 3 pays a total amount of 8 units of money, bidder 4 pays a total amount of 7 units of money, and bidder 5 pays a total amount of $5 + 7 + 8 = 20$ units of money. ▲

Decreasing Auctions. In a decreasing auction the seller starts at a very high price that is gradually decreased. Each bidder calls out a quantity that he is willing to buy at the desired price. The bidder wins this quantity at the current price and the auction terminates when all units have been sold. Bid shading is expected invoking inefficient outcomes. Decreasing auctions are outcome equivalent to sealed-bid discriminatory auctions but not strategically (as in the case of one unit), since information about one's rivals' valuations is being released during a decreasing auction. Under the private values model, this information is useless.

Revenue Comparison among Multi-Unit Auctions. The comparison among the uniform, discriminatory, and generalized Vickrey auctions according to their revenues, does not give a clear rating. If bidding behavior were the same in these formats, then the discriminatory auction would give the highest revenues among the three, the uniform the second highest and the Vickrey the

	Clinching Process				
Price	Bidder 1	Bidder 2	Bidder 3	Bidder 4	Bidder 5
1	-	-	-	-	-
2	-	-	-	-	-
3	-	-	-	-	-
4	-	-	-	-	-
5	-	-	-	-	1
7	-	1	-	1	+1
8	1	+1	1	-	+1
Allocations	1	2	1	1	3
Payments	8	15	8	7	20
Net benefits	2	13	7	5	33

Table 2.4: Ascending Clock Auction with Clinching in Example 2.5.2: clinching process

worst among the three, as can be easily seen in Figure 2.3. But this is not the case, since as we saw previously, bidding behavior differs significantly in each auction. Bid shading for example, may cause a dramatic decrease in revenues in a uniform or a discriminatory auction. Adapting the distribution of bidders' valuations appropriately, one can always find examples in which any of the above auctions may be superior to the others.

2.5.3 Progressive Second-price Auction (PSP): a mechanism for selling an arbitrarily divisible good

Lazar and Semret propose in [19], the Progressive Second Price Auction Mechanism for bandwidth allocation in a single link. They consider a quantity C of an arbitrarily divisible resource, namely bandwidth, to be allocated to a set of bidders $\mathcal{N} = \{1, \dots, N\}$. Each bidder is assumed to have a privately known valuation. Bidder's i valuation is given by the function $\theta_i : [0, C] \to \Re$ and satisfies the following assumptions:

- $\theta_i(0) = 0$,

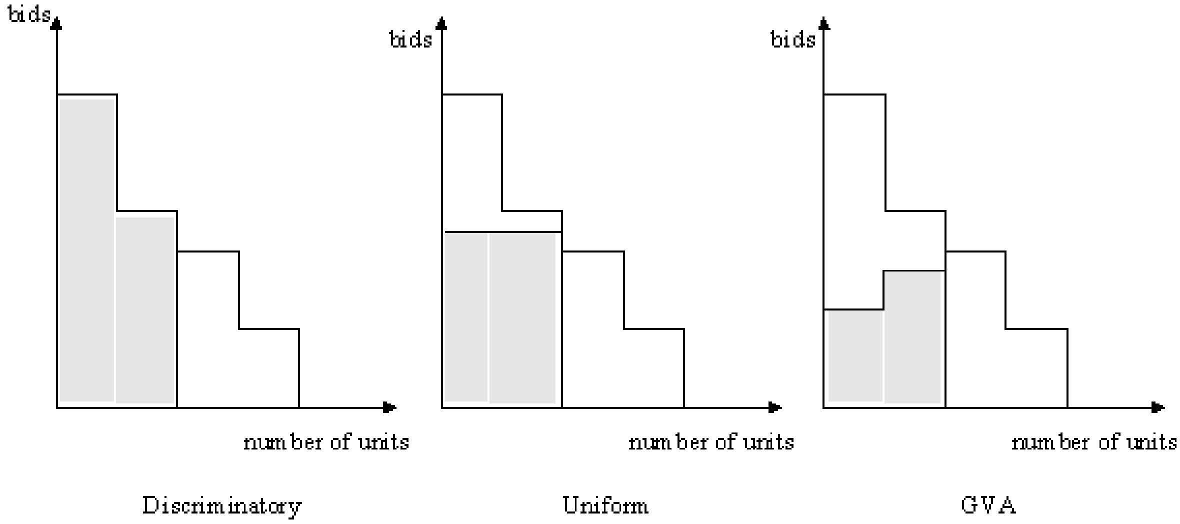

Figure 2.3: Payments under different pricing rules

- θ_i is differentiable,
- $\theta_i' \geq 0$, non-increasing and continuous,
- $\exists \gamma_i > 0, \forall z \geq 0, \theta_i'(z) > 0 \Rightarrow \forall \eta < z, \theta_i'(z) \leq \theta_i'(\eta) - \gamma_i(z - \eta)$. This implies that as long as the valuation is strictly increasing, it must also be strictly concave.

Equivalently, a bidder whose valuation satisfies the above assumptions is characterized to have elastic demand. A repeated game of incomplete information is formed, where each bidder places a bid asynchronously to his opponents, after observing their bids, thus replacing his old bid. New responses from bidders follow, and the procedure terminates when none of them is willing to renew his bid. Bidder's i bid is the pair $s_i = (q_i, p_i)$, where q_i is the desired quantity at a unit price p_i. The allocation is calculated according to the following allocation rule: for a fixed opponent profile (i.e. given the opponents' bids), a bidder gains the minimum of his bid quantity and the maximum available quantity at his bid price. He pays the social opportunity cost; that is the amount offered for this quantity by the bidders who are excluded by his presence. Formally, the allocation and payment rule for bidder i is given by:

$$\text{Allocation: } a_i(s) = \min\{q_i, \underline{Q}_i(p_i, s_{-i})\}, \tag{2.3}$$

$$\text{Payment: } c_i(s) = \sum_{j \neq i} p_j [a_j(0; s_{-i}) - a_j(s_i; s_{-i})], \tag{2.4}$$

where, $s = (s_1, \ldots, s_N)$ is bidders' bid profile, $s_{-i} = (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots, s_N)$ is the bid profile of bidders's i opponents, and [3]

$$\underline{Q}_i(y, s_{-i}) = \left[C - \sum_{p_k \geq y,\ k \neq i} q_k \right]^+, \text{ for } y \geq 0, \tag{2.5}$$

is the maximum available quantity for bidder i for a fixed opponent profile, if those of his opponents that have bidded at or above price y, obtain their demanded quantity. Bidder's i payment is the social opportunity cost in the following sense: for each opponent j of bidder i, the difference $a_j(0; s_{-i}) - a_j(s_i; s_{-i})$ is the quantity bidder j loses by i's presence. For this quantity, bidder i pays the price offered by bidder j, that is $p_j[a_j(0; s_{-i}) - a_j(s_i; s_{-i})]$.

Lazar and Semret analyze the PSP auction as an iterated game of complete information. For a fixed opponent profile s_{-i}, a best reply s_i is defined to be the bid that maximizes bidder's i net benefit. A Nash equilibrium is a bid profile s, if for each i, s_i is i's best reply given s_{-i}. Lazar and Semret introduce the notion of the ϵ-Nash equilibrium. For a fixed opponent profile s_{-i}, they define the ϵ-best reply s_i^ϵ to be the bid that maximizes his net benefit within ϵ: $NB_i(s_i^\epsilon; s_{-i}) \geq NB_i(s_i'; s_{-i}) - \epsilon$, $\forall s_i' \in S_i(s_{-i})$, where $S_i(s_{-i})$ is the set of all feasible bids of bidder i for a fixed opponent profile. An ϵ-Nash equilibrium is a bid profile s^ϵ, if for each i, s_i^ϵ is i's best reply given s_{-i}^ϵ.

Lazar and Semret prove in [19], that under the assumption of elastic demand and for any fixed opponent bid profile, there exists a truthful ϵ-best reply for each bidder, provided that there is always a fixed bid $s_0 = (C, p_0)$ set by the seller. A bidder has to calculate his ϵ-best reply each time he submits a bid. In order to determine his ϵ-best reply, bidder i employs the "staircase" P that is formed by i's opponent bid profile s_{-i}. Then, bidder's i best reply $s_i = (q_i, p_i)$ is the point among those of his marginal valuation function that intersects with the "staircase" P. The ϵ-best reply is then found by reducing quantity by $\epsilon/\theta_i'(0)$, that is, $q_i^\epsilon = [q_i - \epsilon/\theta_i'(0)]^+$, and increasing price such that $p_i^\epsilon = \theta_i'(q_i^\epsilon)$. Due to this equality, the ϵ-best reply is characterized as truthful, in

[3] For any $z \in \Re$, $[z]^+ = z$, if $z > 0$ and $[z]^+ = 0$, if $z \leq 0$.

the sense that it constitutes a point of user's i demand function. The following example illustrates the derivation of the ϵ-best reply.

Example 2.5.3

Consider four bidders that compete for quantity C. The seller has set a reserve price p_0. Assume that bidder 4 wants to submit a bid and observes an opponent profile $s_{-4} = (s_1, s_2, s_3)$, where $s_1 = (q_1, p_1), s_2 = (q_2, p_2)$ and $s_3 = (q_3, p_3)$. Figure 2.4 illustrates the "staircase" P of bidder's 4 opponents. His ϵ-best reply is the point $s_4^\epsilon = (q_4, p_4)$. Had the auction terminated by this moment, bidder 4 would win the quantity q_4 for total charge c_4 equal to the shaded area in Figure 2.4. ▲

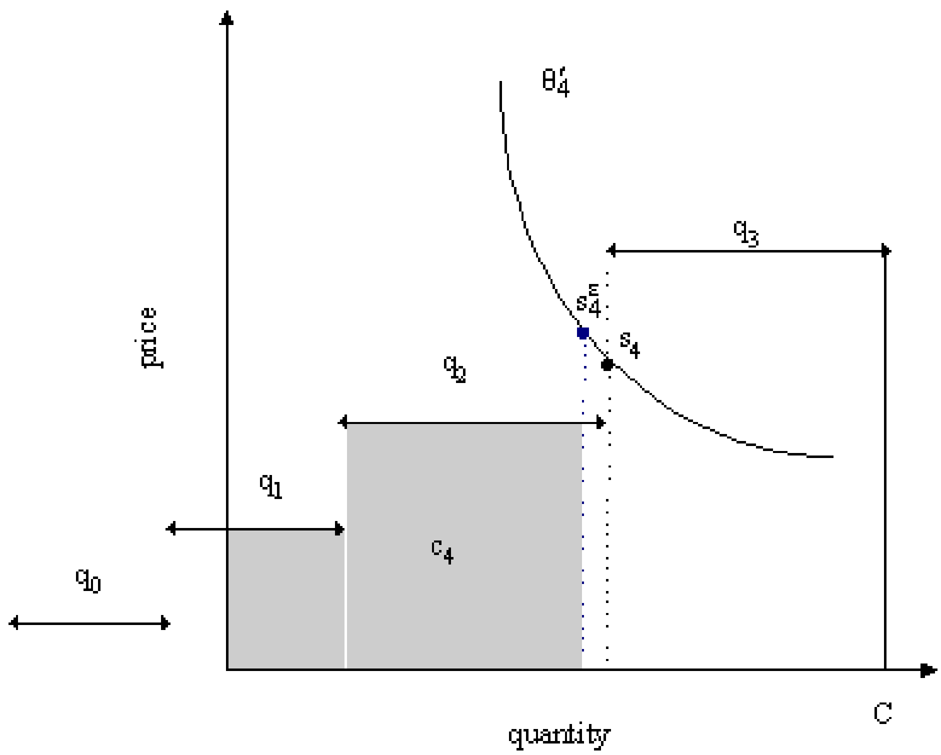

Figure 2.4: Truthful ϵ-best reply for bidder 4

The authors of [19] further prove, that if all bidders follow the strategy of the ϵ-best replies the game will end at an ϵ-Nash equilibrium, for every ϵ. ϵ can be interpreted as a *fee* paid by each bidder each time they submit a bid. In addition, the PSP auction approximates the efficient allocation provided that the second derivatives of bidders' valuation functions are bounded. In particular, the social welfare of the outcome is very close to its optimal value. The difference of the two values is bounded by a term that depends on the value of ϵ and on the bound of second derivative of bidders' valuation functions. The smaller ϵ, the closer the outcome to the efficient allocation gets, and the longer the procedure takes to converge. The authors further extend their work to the network case. We examine thoroughly the network-wide PSP in Chapter 5, where we introduce and assess an improvement.

2.6 Multi-Object Auctions

Below we briefly discuss the main issues on multi-object auctions for completeness reasons. Such are not considered in this dissertation. Auctions for heterogeneous goods are further classified as *combinatorial auctions* or *independent bidding auctions*. In combinatorial auctions, bidders are allowed to submit a single bid for combinations of goods. Winner determination differs from the way winners are determined in the multi-unit auctions. In a multi-object auction winners are the bidders whose sum of bids is the maximum among all feasible combinations of bids. A feasible combination of bids is a set of bids that includes each good exactly once. Combinatorial auctions are appropriate when complementary goods are being sold. Consider for example, the case where bidders compete for bandwidth in each one's path. In this case a combination bid is a bid for bandwidth for all links that form one's path. If combinations of bids are not allowed, then the bidder faces the risk of winning different portions of bandwidth (or no bandwidth at all) in some links. The allocation then would not be efficient, while a portion of the allocated bandwidth would be unused. Combinatorial auctions give rise to the *threshold problem* in which, groups of low-valuation bidders may not be able to coordinate their bids to displace a large (meaning having high valuation) but otherwise inefficient bidder, thus resulting in an inefficient outcome. The following example illustrates the threshold problem.

Example 2.6.1

Consider three bidders competing for goods A and B. A simultaneous ascending auction is performed. Each bidder's valuation is privately known to him according to the following: the valuation of bidder 1 for good A is $\theta_{1,A} = 7$, the valuation of bidder 2 for good B is $\theta_{2,B} = 8$ and valuation of bidder 3 for the combination of goods A and B is $\theta_{3,AB} = 12$. The efficient outcome is to give bidder 1 the good A and bidder 2 the good B (since social welfare is then maximized: $7 + 8 > 12$). However, assume that the auction has come to the following situation: bidder 1 has placed a bid of 1, bidder 2 has placed a bid of 2 and bidder 3 has placed a bid of 11. Neither of the single bidders 1 and 2 can unilaterally place a bid of 11 or 10 respectively so as to overbid bidder 3. Thus, the good will be awarded to bidder 3 which is not the efficient outcome. ▲

Furthermore, in combinatorial auctions a computational problem arises. The number of combinations of goods upon which bids are allowed to be placed is exponential in the number of goods: in an auction where C goods are being sold, the different acceptable combinations of goods equal $2^C - 1$. Thus, the number of submitted bids can be very large, making the problem of winner determination an NP-complete one. This computational complexity can be smoothed by restricting the allowable combinations. For example, in the case of a linear communication network only combinations of adjacent links are meaningful. This problem is considered in [29] and [30].

Independent bidding auctions perform better if supplementary goods are being sold. On the contrary, if complementary goods are being sold, independent bidding induces the *exposure problem*, which is, in a way, opposite to the threshold problem. This is briefly as follows: A bidder is exposed to the risk of not acquiring the desired bundle of goods but a part of it at a total price exceeding this part's value for him, or may not be able to compete for the bundle through each good separately due to the low value each single good has. This may result in bidder losses and/or inefficient outcomes. The following example taken from [8], illustrates the exposure problem.

Example 2.6.2

Consider two bidders competing for goods A and B. A simultaneous ascending auction is performed where only independent bids for the goods are allowed. Prices for both goods increase continuously and each bidder informs if he is active (participates) or not. Each bidder's valuation is privately known to him and is given in Table 2.5. (Goods A and B are complements for bidder 1, while substitutes for bidder 2.) The efficient outcome is to give bidder 1 both goods A and B. However, assume that the auction has come to the following situation: prices for goods A and B have been increased to 1 and 2 respectively. If bidder 1 withdraws from both goods, the auction terminates and the outcome is inefficient: bidder 2 acquires both goods. If bidder 1 continues being active

bidder	θ_{A}	θ_{B}	θ_{AB}
1	1	2	6
2	3	4	5

Table 2.5: Marginal valuations in Example 2.6.2

in higher prices wishing to acquire both goods, then he will win at most one good [4] at a price higher than his own valuation on that good, incurring a loss for him. ▲

Vickrey-Clark-Groves Mechanism (VCG). Consider a set $\mathcal{C}$ of C distinct (not identical) goods that are to be sold to a set of N bidders. Each bidder has a privately known valuation for each combination of the goods (the number of different combinations is 2^{C-1}): for any $K \subset \mathcal{C}$, let $\theta(K)$ denote the total value derived from the K goods. The VCG mechanism introduced by Groves in [11], is a generalization of the GVA for identical units as well as the pivotal mechanism for public goods introduced by Clarke in [7]. The GVA mechanism allows combinatorial bidding. The allocation and payment rules are as follows: winning bids are those non-overlapping bids that maximize social welfare. Each winner pays the difference between the welfare (sum of bids) of the others without his participation and the welfare of the others with him participating. Equivalently, each winner pays his bid reduced by the increment of welfare caused by his participation. Note, that in the case of a single unit, VCG reduces to the Vickrey auction, while in the case of identical units it coincides with the GVA. Truthful bidding is a weakly dominant strategy and thus, VCG is an *incentive compatible and efficient mechanism.* An even more remarkable property of VCG that is due to Krishna and Perry ([17]), is that VCG results in the highest revenues among all efficient mechanisms. The following example illustrates the VCG mechanism.

Example 2.6.3

Consider three bidders competing for two goods A and B. The VCG mechanism is performed. Each bidder's valuation is privately known to him and is given in Table 2.6. (Goods A and B are complements for bidder 3, while substitutes for bidders 1 and 2.)

[4] Bidder 1 wins at most one good, since there are no prices p_1 and p_2 such that $p_1 + p_2 \leq 6$ and $p_1 > 3$ and $p_2 > 4$.

bidder	θ_{A}	θ_{B}	θ_{AB}
1	10	0	10
2	0	15	15
3	0	0	20

Table 2.6: Marginal valuations in Example 2.6.3

The outcome is to allocate good A to bidder 1 and good B to bidder 2 ($10 + 15 > 20$), which is the efficient one. Bidder 1 pays the amount $c_1 = 20 - 15 = 10 - (25 - 20) = 5$: if bidder 1 were not present, bidder 3 would win both units and the winners' social welfare would be 20. If bidder 1 participates, the winners', apart from him, social welfare is 15. Thus the cost for him is the difference $20 - 15$. Similarly, bidder 2 pays the amount $c_2 = 20 - 10 = 15 - (25 - 20) = 10$. ▲

Chapter 3

An Efficient Auction Mechanism for Hierarchically Structured Bandwidth Markets

3.1 Introduction

Nowadays, markets of services and resources involve chains of multiple providers. Thus, in the general case of a business model, there may be multiple levels of bandwidth trading until the resources are actually allocated to the end customers. Indeed, in the presence of large and geographically distributed sets of customers, their direct transaction with the seller is either impossible, or entails high computational and physical or communication overheads. This motivates studying a hierarchical market model. Indeed consider a seller that faces a large and widespread market for selling bandwidth. It is possibly beneficial and more effective for him to assign the trading of bandwidth to *intermediaries*, rather than dealing directly with customers. The fact that intermediaries may have access to a wide number of customers, leads the bandwidth market to grow in size substantially. Moreover, the division of the market into several parts facilitates the selling process, while reduces and/or distributes the management overhead. However, the seller should ensure that the intermediaries' profits, which correspond to a potential loss for him, will be lower than the cost induced by selling bandwidth directly. Furthermore, bandwidth markets are not stable due to unpredictable demand and increasing supply. Traders may acquire excess bandwidth that cannot be stored for later use. The inefficiencies that arise due to the failure of the market to balance, make reselling unavoidable. Profits can be extracted from successive trades, thus help wholesale providers survive and grow in the market. Considering hierarchically structured bandwidth markets, the seller takes into account resale imposing rules that direct the outcome to serve his own purposes. In the degenerate case, the hierarchy consists of only one level where no resale appears. In a middle case some intermediaries proceed on reselling and others keep all the bandwidth for themselves. Another reason that motivates the use of multiple levels in selling bandwidth, is the need to offer differentiated services to the customers. Service provisioning is built upon bandwidth trade giving rise to the necessity of studying models that describe value chains. Note also that the objective of the top-level seller may be either to maximize his expected profit or to allocate the bandwidth efficiently to the end customers.

In this chapter, we formulate a *new* resource allocation problem. In particular, we deal with a two-level hierarchical business model for selling bandwidth in a single link. In the top level the social planner sells bandwidth to intermediate providers [e.g. Internet Service Providers (ISPs)], who in turn sell their assigned shares of bandwidth to the customers in the lower level. The social planner defines the rules of the mechanism. Our objective is the *efficient overall allocation* of the entire supply of bandwidth to the customers, as if the social planner were to assign this bandwidth directly. Hence the use of the term social planner, rather than profit-seeking seller.

Due to its hierarchical structure, the above allocation problem cannot be solved efficiently in two independent stages, one for each level. The social planner cannot sell the bandwidth to the providers before they acquire knowledge of the demand they will face by their customers, while the providers cannot trade with their customers before they learn the quantity of bandwidth they obtain from the social planner. Thus, a dynamic mechanism enforcing coordination between the two levels is required. The exchange of information between the auctions taking place at the two levels is necessary for the solution to be efficient. In light of this necessity, we propose an innovative mechanism comprising an appropriate auction in each level. The auctions are coordinated so that supply is exhausted at the end, as required for attaining efficiency. The service providers and the customers are expected to act according to their *own incentives*, i.e. so that their respective benefits are maximized, without being concerned about social welfare. Nevertheless, our mechanism is specified so that bidders of both auctions have the incentives both to participate uninhibitively and to *bid truthfully* (incentive compatibility property) and efficiency be ultimately attained. Our mechanism has the following important feature: Although bidding of each provider is based on aggregation of information of his own market, there is no loss of information regarding the ordering of customers' valuations over all markets. Such a loss of information would clearly lead to an inefficient outcome.

What is of particular importance is that efficiency and incentive compatibility are influenced by the way *information* is distributed to the players. Efficiency and incentive compatibility are attained in an environment where each player only possesses a certain part of information on the entire market. In particular, we assume that customers' preferences are privately known to them, providers have no prior information about their local markets, while the social planner may only interact with providers. Moreover, the rules of the mechanism are announced to all players by the social planner but the bidding records of both levels are only released after the end of the procedure. The social planner restricts providers: a) to supply their local markets with bandwidth as soon as it is made available to them, and b) to sell each unit of bandwidth at a price that does not exceed the one of the most expensive unit sold to any of the providers at the top-level auction. These rules render our mechanism efficient.

Further, we formulate business models in which the trading rules are not imposed exclusively by the social planner. In particular, we extend the model to a less restrictive one, in which providers are allowed to choose any payment rule to apply to their own local market, provided they conform to certain (yet fewer) restrictions. We investigate cases in which it is possible for providers to gain more profits by choosing payment rules that lead to inefficient allocations. However, we prove that if each customer is allowed to *select* his own provider on the basis of the selected payment rules, then each provider has the incentive to *keep* the original rule, for otherwise he would end up with no customers.

Walsh and Wellman [34] consider task allocation problems with hierarchical dependencies, however in a different context. They consider a supply chain through a hierarchy of task achievement; a supplier supplies a set of primitive goods, a provider uses primitive goods as input to produce a single output good and finally the consumer acquires a high-level good. Each consumer wishes to acquire only one high-level good. In that paper the term "hierarchy" refers to the various steps taken by different providers to produce and dispose a good to the customer rather than to the process of selling goods through a set of intermediaries. In the absence of providers high-level goods cannot be produced and sold by the supplier, whereas in the absence of intermediaries in our problem the social planner sells the goods directly to the customers.

Maillé and Tuffin [24] propose a charging scheme applied in a backbone network interconnected with congested access networks that have a tree structure. The authors perform a multi-bid auction to allocate bandwidth in the backbone network and then assign the bandwidth in the same way in the access networks. The hierarchical structure involves the topology of the network presented in [24] rather than the selling procedure as we consider it in our problem.

The remainder of this chapter is organized as follows: We formulate the hierarchical bandwidth allocation problem, and analyze the requirements for an efficient solution to the problem in Section 3.2. In Section 3.2.2 we discuss various mechanisms for the case of a single unit of bandwidth. In Section 3.3 we develop our mechanism to sell bandwidth hierarchically. We then analyze players' strategies and prove that the proposed mechanism is incentive compatible and efficient in Section 3.4. In Section 3.5 we extend the model to allow providers determine by themselves the payment rules applied to their own markets, and study the related consequences. Finally, in Section 3.6 we provide some concluding remarks.

3.2 The Hierarchical Bandwidth Allocation Problem

3.2.1 Problem Specification

We consider the problem of allocating C indivisible units of bandwidth in a single network link efficiently. Allocation is performed in two levels (C is taken as integer). We identify three types of players: the *social planner* who sells the bandwidth to a set of M bandwidth *providers* in the top level, who in turn sell the obtained quantity of bandwidth to a set of N $(N > M)$ *customers* in the lower level. The social planner imposes certain allocation and payment rules for both levels.

We assume that the quantity C of bandwidth available in the top level is fixed and known to all players, and that customers are partitioned into *fixed* groups, each of which constitutes the local market of a certain provider. We denote the local market of provider j as the set of customers S_j, for $j = 1, \ldots, M$. Without loss of generality we assume that the indices of customers within the same market j constitute a subset of $|S_j|$ consecutive numbers of $\{1, \ldots, N\}$ that is,

$S_j = \{|S_{j-1}| + 1, \ldots, |S_{j-1}| + |S_j|\}$. Note that customers are *not* allowed to move to another market after their initial choices. This issue is revisited in Section 3.5. Each customer i, receives marginal utility $\theta_{i,k}$ by his k^{th} allocated unit of bandwidth, for $k = 1, \ldots, C$. We assume that marginal utilities of each customer are diminishing and privately known to him. A customer knows neither the utility function of other customers, nor the distribution this utility is drawn from, not even the quantity to be offered in his local market. In contrast to customers, providers do *not* know their own valuations. Indeed, provider's j marginal utility $u_{j,k}$ is assumed to equal the revenue he would obtain if he were to sell the k^{th} unit of bandwidth after the trade. Thus, it cannot be determined or predicted without knowledge of market demand, which is taken to be completely unknown. As a consequence, all parts have incomplete demand information that must be extended through the mechanism process. Each part possesses a certain piece of information and is lacking of others. In particular, customers do not know the quantity to be offered in their local market, as it is determined by the competition among the providers. At the same time, providers need to exploit their local demand (through the competition among the customers), in order to obtain the optimal bandwidth quantity for their own local markets by participating in the top-level auctioning mechanism.

The objective of the social planner is to *maximize social welfare*, which measures the overall well-being of the society. In our setting it is given by the formula:

$$\text{SW}(\vec{x}) = \sum_{i=1}^{N} \sum_{k=1}^{x_i} \theta_{i,k}, \tag{3.1}$$

where $\vec{x}$ is the vector of the allocated quantities to the N customers. Note that, without loss of generality, we assume that the cost of bandwidth to the social planner is zero, since this cost is already sunk. The social planner wishes to determine the vector $\vec{x}^*$ of the allocated quantities that maximizes the social welfare given the bandwidth constraint. In case of complete information, the social planner would solve the problem by ordering the $\theta_{i,k}s$ over all customers and selecting the C largest ones.

From the social planner's perspective, the mechanisms of the two levels need to appropriately interact so that both providers and customers dynamically obtain the exact set of information that leads to an efficient outcome. It is necessary that we only employ efficient mechanisms in both levels. Still, this is only a necessary condition for efficiency in the context of our problem, but it is *not* a sufficient one. To illustrate this we will give an example in the case of a single unit, where coordination between the two levels fails, even though efficient mechanisms for both levels are employed.

3.2.2 The case of a single unit of bandwidth

Analyzing the case of selling a single unit of bandwidth provides us with considerable insight and understanding on the issues that arise in the general case $C > 1$ too. In particular, we discuss about various well-known auction mechanisms for a single unit and whether they can be applied to our problem of hierarchical allocation.

Suppose that the social planner wants to sell a single good (i.e., unit of bandwidth) to one of N customers with valuations θ_i through M providers. He performs an auction to select the winner among the providers (top-level auction), who will sell the good to one of his customers by means of another auction (lower-level auctions). The fact that the good must be sold to a provider and subsequently to one of his customers does not necessarily determine the order the two auctions take place, although it imposes certain limitations thereto. The objective of the social planner is to ultimately sell the good to the customer with the highest valuation. Below we examine three mechanisms, only one of which attains efficiency under the assumption of privately known valuations.

First we deal with the combination of Vickrey auctions in both levels. The lower-level auction in each market is performed first so that the respective provider learns his utility. That is, his revenue if he does win the good in the top-level auction; this revenue equals the second highest bid among the bids placed by his customers. Then, the providers participate to the top-level auction according to this utility. All players (in both levels) bid truthfully, due to the Vickrey payment rule, but the outcome may *not* be efficient. That is, the good may not be sold to the provider with whom is associated the customer having the overall highest valuation. Indeed, the good is sold to the provider who submitted the largest among the second highest valuations per market, which does not necessarily coincide with the aforementioned provider. The same results hold in case of applying the English auction independently in the two levels starting from the lower one: each provider learns his potential revenue which is again the second highest valuation of his customers.

In the second mechanism under consideration, the First Price sealed bid auction is applied first in the lower level and then the Vickrey auction in the top level. Since the payment rule (in the lower level) is pay-your-bid in case of winning, the customers shade their bids. For each provider, the highest of these bids is his potential revenue. Thus, it is used as his bid in the top level, since he bids sincerely due to the Vickrey payment rule. However, it is possible that the order of providers according to the highest shaded bid per local market is different than that according to the highest valuation, thus resulting in inefficiency.

Suppose now that we apply the English auction in both levels simultaneously as follows: a common price ascends continuously in each level. At every price, each customer either accepts (e.g. by pressing a button) the offer according to their valuations, or withdraws from the auction. Regarding each provider's strategy, it is meaningful to accept the offer at the same price if at least one of his customers accepts it; otherwise, the provider withdraws from the top-level auction. This auction terminates at the first price where only one of the providers still accepts the offer. This is the winner and pays the current price. His market's auction continues until only one of his customers still accepts the offer. This customer is the final winner and pays the final price of his local market. The winning provider surely ends up with non-negative profits, because his buy price is always less than or equal his sell price. The mechanism produces the efficient outcome, since as the price increases only the customer with the highest valuation remains active. The provider's strategy described above is a weakly dominant strategy: if he withdraws instead of accepting the offer at a given price where some of his customers are still active, then he ends up with a zero profit, (because he will not win anything); if he accepts instead of withdrawing (when there is no more activity in his local market) then he may end up with a negative profit (lower selling price than buying if he wins in the top level), or zero profit.

Example 3.2.1 clarifies the above ideas.

Example 3.2.1

Assume two providers A and B having two and three customers with valuations $\theta_1 = 2,\ \theta_2 = 7,\ \theta_3 = 4.5,\ \theta_4 = 6,\ \theta_5 = 3$ respectively. For the outcome to be efficient, the good must be sold to customer 2.

a. We apply the first mechanism (i.e. Vickrey auctions in both levels) with privately known valuations. The bids in the two levels are shown in Table 3.1. Providers A and B bid their revenues

	Lower Level A		**Lower Level B**			*Top Level*	
	Cust1	Cust2	Cust3	Cust4	Cust5	ProvA	ProvB
bid in mechanism 1	2	7	4.5	6	3	2	4.5
bid in mechanism 2	1	3.5	3	4	2	3.5	4

Table 3.1: Bidding process in Example 3.2.1 a and b

from their customers in case of winning; that is, 2 and 4.5 respectively. Provider B wins the good at price 2 and sells it to customer 4 at price 4.5. Thus, the efficient outcome is not attained.
b. Next, we apply the second mechanism assuming that the customers are symmetric and their valuations have been drawn from the uniform distribution on $[0, 10]$, which is common knowledge to all players. In this case the equilibrium strategy of customers is to perform bid-shading according to the size of the respective local market (see [35]). Thus, customer i would bid $\beta_i = \frac{N_j - 1}{N_j}\theta_i$, for all $i = 1, 2, 3, 4$ and $j = 1, 2$. The bids of all players are shown in Table 3.1. Provider B wins at price 3.5 and sells the good to customer 4 at price 4. The outcome is not efficient because the customers's shading does not preserve the ordering of their valuations since $\beta_2 < \beta_4$ although $\theta_2 > \theta_4$.
c. The third mechanism assumes privately known valuations for the customers. The price starts at 1 and at each price each player accepts the offer ("yes") or withdraws ("no") as shown in Table 3.2. Note that although the price is continuous we only present the price values at which new withdrawals arise. Provider A wins at price 6 and sells the good to customer 2 at price 6, which is the efficient outcome. ▲

The above discussion supports our assertion that our problem cannot be solved without the appropriate relation of information in the two levels. The third mechanism attains efficiency, because the order of customers' valuation is preserved in the top-level auction in contrast to the other mechanisms that just exploit a certain point of the local demand. In the generalized problem of selling C units, the Ascending Clock Auction with Clinching (ACC) proposed by Ausubel [1], which is the generalization of the English Auction, is not a proper solution, as explained in the next section. On the contrary, a new version of ACC with a modified payment rule serves our purpose.

	Lower Level A		**Lower Level B**			*Top Level*	
Price	Cust1	Cust2	Cust3	Cust4	Cust5	ProvA	ProvB
1	yes	yes	yes	yes	yes	yes	yes
2	no	yes	yes	yes	yes	yes	yes
3		yes	yes	yes	no	yes	yes
4.5		yes	no	yes		yes	yes
6		yes		no		yes	no

Table 3.2: Bidding process in Example 3.2.1 c

3.2.3 Providers' profits

Employing efficient mechanisms in both levels of our hierarchy is indeed not sufficient to achieve overall efficiency as we showed in Example 3.2.1. In the next subsection we prove that a necessary condition for overall efficiency is that each provider submits *truthfully* his market demand in the top level auction. Thus, the social planner has to design the whole mechanism so that each provider maximizes his profits by transferring his market demand in the top level. A requirement for this, is that all players obtain non-negative profits at the end. Negative profits raise participation issues and generate incentives that result in inefficiencies. Since customers know their own valuations and act rationally, there is no possibility of obtaining negative profits in any mechanism. But this is not the case for the providers, since they participate in two different trades. Thus, depending on the mechanism, there may exist a possibility of price inconsistency for a provider; i.e. to buy certain units of bandwidth at prices that are higher than the corresponding selling prices. In Example 3.2.1 a. in which Vickrey auctions are employed in both levels, provider's A local market demand is one unit at price 7. If provider A submits truthfully his markets' demand or equivalently if he submits a bid equal to 7, he will obtain the good at price 4.5 but he will sell it to customer 2 at a price of 2, thus making a negative profit equal to $2 - 4.5 = -2.5$.

3.2.4 The Hierarchical Optimization Problem

Next, we define an optimization problem that is motivated by the hierarchical nature of our model, and prove that it produces the optimal solution $\vec{x}^*$ that maximizes social welfare. Note that we specify the appropriate optimization problem that serves our purposes. The solution will be given in a subsequent section.

HIERARCHICAL OPTIMIZATION:

LOWER LEVEL ALLOCATION

$$\begin{aligned} &\max_{x_i,\ i\in S_1} \sum_{i\in S_1}\sum_{k=1}^{x_i}\theta_{i,k} & & & &\max_{x_i,\ i\in S_M} \sum_{i\in S_M}\sum_{k=1}^{x_i}\theta_{i,k} \\ &\text{s.t.} \sum_{i\in S_1} x_i = q_1^* & &\cdots & &\text{s.t.} \sum_{i\in S_M} x_i = q_M^* \\ &x_i \in \aleph,\ \forall\, i \in S_1 & & & &x_i \in \aleph,\ \forall\, i \in S_M \end{aligned} \tag{3.2}$$

where $\vec{q}^* = (q_1^*, \ldots, q_M^*)$ is the vector of the allocated quantities to the M providers and is the solution of:

TOP LEVEL ALLOCATION

$$\begin{aligned} &\max_{q_1,\cdots,q_M} \sum_{j=1}^{M}\sum_{k=1}^{q_j} v_{j,k} \\ &\text{s.t.} \sum_{j=1}^{M} q_j = C \\ &q_j \in \aleph,\ j = 1, \cdots, M. \end{aligned} \tag{3.3}$$

The values $\{v_{j,k} : k = 1, \cdots, C\}$, are the marginal utilities of market j, which are defined to be the C maximum marginal utilities over all customers of market j. That is:

$$\begin{aligned} &v_{j,k} = k^{\text{th}} \text{ maximum of } \{\theta_{|S_{j-1}|,1}, \cdots, \theta_{|S_{j-1}|,C}, \cdots, \theta_{|S_{j-1}|+|S_j|,1}, \cdots, \theta_{|S_j|,C}\}, \\ &\text{for } k = 1, \cdots, C \text{ and } j = 1, \cdots, M. \end{aligned} \tag{3.4}$$

The interpretation of the above optimization problem is as follows: the social planner supplies the provider of each local market j, whose utility is v_j, with the optimal, in terms of efficiency, quantity q_j^* (top level of hierarchy). Provider j in turn, allocates q_j^* optimally, in terms of efficiency, to its set S_j of customers (lower level of hierarchy). We claim that this procedure provides the overall efficient allocation $\vec{x}^*$. The proof of the following proposition is given in Appendix A.

Proposition 3.2.1 *If the marginal utilities $v_{j,k}$ for $j = 1, \cdots, M$ and $k = 1, \cdots, C$ are given by Equation (3.4), then the hierarchical optimization problem is equivalent to the social welfare maximization given by Equation (3.1).*

Yet, we have assumed that the social planner has no knowledge of the various θ_is and therefore cannot derive the q_j^*s. Therefore, the social planner should run an auction in order to acquire the necessary information from the providers. However, the social planner should set the stage so that the allocation of bandwidth to providers and thereby to their local markets does induce such incentives so that the hierarchical optimization problem is solved. Specifically, Proposition 3.2.1 implies that the social planner, in order to allocate bandwidth efficiently to the N customers, has to define a dynamic procedure in which:

1. Customers reveal their true marginal valuations (θ_i), so that demand for each local market j is derived by the values $v_{j,k}$, for $k = 1, \cdots, C$.
2. Providers reveal their local market demand to the social planner so that ensuing efficient allocation of bandwidth coincides with $q_1^*, \ldots, q_M^*$.
3. Each provider allocates q_j^* efficiently to his customers.

The mechanism we propose in Section 3.3 satisfies all these requirements.

3.3 The Hierarchical Auction Mechanism

We propose a synchronized [1] mechanism in the top and lower levels of our hierarchy for selling C indivisible units of bandwidth to N customers through M providers. The Ascending Clock Auction with Clinching is performed in both levels. In the lower level, however, we introduce a *new* allocation rule, which makes use of the outcome of the top-level auction; this rule is discussed below. The price is common in the two auctions, starts at a reserve price p_0 and increases

[1] A non-synchronized version is presented in 3.4.3.

continuously until the end of the process at a final price, to be defined below. Without loss of generality, we henceforth assume that $p_0 = 0$.

Players of both levels bid for quantities at every price according to their strategies. Similarly to the zero information model of ACC (see Section 2.5.2), we assume that during the bidding process *no* information is made available to any player about his opponents' bids. Each customer observes price p and decides whether to submit a new bid at this price or not. A pure strategy for customer i is the function $x_i : \Re \rightarrow \aleph$, where $x_i(p)$ denotes the quantity demanded at price p. Each customer's bids have to be non-increasing as price ascends. That is, $x_i(p') \leq x_i(p)$ for any two prices p and p' where $p \leq p'$. Moreover, $x_i(p)$ is right-continuous. Each provider observes his own customers' demanded quantities at price p and uses this information to calculate his bid for the top-level auction, as we discuss in Section 3.4. This bid is submitted automatically in the top-level auction. (Practical matters are discussed in subsection 3.4.3.) A pure strategy for provider j is the function $Q_j : \Re^{N_j} \rightarrow \aleph$, where N_j is the number of customers in market j. $Q_j(\vec{x}(p))$ denotes the quantity provider j demands in the top-level auction at price p. For each price p, this depends on the vector $\vec{x}(p) = (x_1(p), \cdots, x_{N_j}(p))$ of demanded quantities in the provider's local market. In order to simplify notation, we denote this function as $Q_j(p)$. Again, the bids of provider j bids have to be non-increasing as price ascends. We will determine the dominant strategies of all players in the next section. The top-level auction terminates at the first time where demand equals supply in the top-level auction. Let $p_{\rm f}$ be the final price of the *top-level* auction. Accordingly, each lower-level auction terminates when demand equals supply at the respective market.

The evolution of the mechanism is given completely by the set of prices p^l, $l = 1, \cdots, L$, corresponding to the *occasions* on which one or more players *in any of the two levels* strictly decreases his quantity. Next, we will define the allocation and payment rules for the two levels, using the notation of [1].

TOP LEVEL AUCTION:

Let q_j^l denote the quantity demanded by provider j at the $l^{\rm th}$ occasion. Due to clinching in the top-level auction, at each price p^l, each provider has already guaranteed a quantity of bandwidth for his respective market. We define the quantity c_j^l clinched up to (and including) price p^l by provider j as follows:

$$c_j^l = \max\{0, C - \sum_{k \neq j} q_k^l\}, \text{ for } l = 1, \cdots, L \text{ and } j = 1, \cdots, M. \quad (3.5)$$

After the bidding process is completed, the social planner announces to each provider the quantity of bandwidth he wins and his charge. In particular, each provider obtains the quantity he demanded at the final price p_{f} and pays for each unit of bandwidth the standing price at which he clinched this unit, as suggested by the Ascending Clock Auction with Clinching. Formally, the outcome of the top-level auction is defined by:

$$\text{Allocation of Provider } j: \; q_j^* = c_j^L = q_j^L, \text{ for } j = 1, \cdots, M; \quad (3.6)$$

$$\text{Payment of Provider } j: \; k_j(q_j^*) = \sum_{l=1}^{L} p^l \cdot (c_j^l - c_j^{l-1}), \text{ for } j = 1, \cdots, M. \quad (3.7)$$

We refer to the case where one provider obtains the entire capacity as the case of *non-competitive market*. The typical case in which two or more than two providers share the capacity is referred to as the case of *competitive market*.

LOWER LEVEL AUCTION:

Let x_i^l denote the quantity demanded by customer i at the l^{th} occasion. That is, $x_i^l = x_i(p^l)$. For the allocation and payment by the customers (in the lower level) we introduce a *new clinching rule*: At each price p, and at each local market of provider j, the condition for determining the quantity to be clinched by the various customers employs the already guaranteed supply c_j^l in this market. That is, as long as a provider clinches new units of bandwidth in the top-level auction, he is required to make them available in the lower auction of his own market. This is different than performing the original ACC auction for the c_j^L units of bandwidth the provider wins at the end of the top-level auction. We denote as b_i^l the quantity clinched by customer i up to (and including) price p^l; this is given by:

$$b_i^l = \max\{0, c_j^l - \sum_{k \neq j} x_k^l\}, \text{ for } l = 1, \cdots, L \text{ and } i = 1, \cdots, N, \quad (3.8)$$

where c_j^l is the supply offered at the l^{th} occasion at the customer's i local market j. Each customer wins the quantity demanded at the final price p_{f} and is charged according to two restrictions, which are part of the definition of our mechanism:

1. Each customer pays the standing prices at which he clinched the won units of bandwidth as defined in the modified process above.

2. If the case of competitive market applies, then the following rule is also imposed: no unit of bandwidth is sold at a price higher than p_{f}.

Formally, the outcome of the lower-level auction is defined by:

$$\text{Allocation of Customer } i: \; x_i^* = b_i^L = x_i^L, \; i = 1, \cdots, N, \tag{3.9}$$

Payment of Customer i in the case of competitive market:

$$w_i(x_i^*) = \sum_{l=1}^{L} (\min\{p^l, p_{\mathrm{f}}\}) \cdot (b_i^l - b_i^{l-1}), \; i = 1, \cdots, N$$

Payment of Customer i in the case of non-competitive market:

$$w_i(x_i^*) = \sum_{l=1}^{L} p^l \cdot (b_i^l - b_i^{l-1}), \; i = 1, \cdots, N. \tag{3.10}$$

In Example 3.3.1, we apply the proposed mechanism assuming that players employ the following strategies: a) each customer bids truthfully according to his utility function and b) each provider bids the aggregate demand of his local market observed at each price provided that it is less than the capacity C. If the aggregate demand exceeds supply, the provider submits the capacity C. In Section 3.4 we will prove that these are indeed the dominant strategies for the customers and the providers respectively.

Example 3.3.1

Suppose there is an amount of $C = 8$ units of bandwidth that is allocated to two service providers A and B that have two customers and three customers respectively, with marginal valuations shown in Table 3.3. The efficient outcome is given by the vector $\vec{x}^* = (x_1^*, x_2^*, x_3^*, x_4^*, x_5^*) = (1, 2, 1, 1, 3)$ and the optimal social welfare is $SW(\vec{x}^*) = 10 + 17 + 11 + 15 + 12 + 19 + 18 + 16 = 118$. We apply the proposed procedure letting the price start at price 1. We provide two tables to describe the evolution of the hierarchical auction. Table 3.4 presents the bids for both levels at the prices of various occasions. Note that the bids of the providers are provided in the last two columns. This property is proved in Section 3.4. All auctions terminate simultaneously. The last row shows the final allocation which is the efficient allocation $\vec{x}^*$. After the procedure is terminated, the social

Provider A		**Provider B**		
Cust1	Cust2	Cust3	Cust4	Cust5
$\boldsymbol{\theta_{1,1} = 10}$	$\boldsymbol{\theta_{2,1} = 17}$	$\boldsymbol{\theta_{3,1} = 15}$	$\boldsymbol{\theta_{4,1} = 12}$	$\boldsymbol{\theta_{5,1} = 19}$
$\theta_{1,2} = 4$	$\boldsymbol{\theta_{2,2} = 11}$	$\theta_{3,2} = 7$	$\theta_{4,2} = 8$	$\boldsymbol{\theta_{5,2} = 18}$
$\theta_{1,3} = 2$	$\theta_{2,3} = 1$	$\theta_{3,3} = 5$	$\theta_{4,3} = 3$	$\boldsymbol{\theta_{5,3} = 16}$

Table 3.3: Marginal valuations in Example 3.3.1

	Lower Level A		**Lower Level B**			*Top Level*	
Price	Cust1	Cust2	Cust3	Cust4	Cust5	ProvA	ProvB
1	3	2	3	3	3	5	8
2	2	2	3	3	3	4	8
3	2	2	3	2	3	4	8
4	1	2	3	2	3	3	8
5	1	2	2	2	3	3	7
7	1	2	1	2	3	3	6
8	1	2	1	1	3	3	5

Table 3.4: Bidding process in Examlpe 3.3.1

planner and the providers calculate their allocations and payments, as shown in Table 3.5. Note that the providers' payments are given in the columns 2 and 3 prior to the customers' payments. Customer 2 for example, clinches his first unit of bandwidth at price 7 since at this price the total supply for provider A is 2 units and his opponent's (customer 1) bid is 1 unit, thus the residual supply left for customer 2 equals 1 ($= 2 - 1$). In the sequel, at price 8 the supply for provider A equals 3 units, meaning that customer 2 clinches his second unit of bandwidth at this price. ▲

3.4 Derivation of Players' Strategies and Efficiency

In this section, we analyze players' strategies and prove that all the objectives set by the social planner are met when the proposed allocation mechanism is applied. We will prove that in our

	Top level		**Lower level A**		**Lower Level B**		
Price	ProvA	ProvB	Cust1	Cust2	Cust3	Cust4	Cust5
1	-	3	-	-	-	-	-
2	-	+1	-	-	-	-	-
3	-	-	-	-	-	-	-
4	-	+1	-	-	-	-	-
5	1	-	-	-	-	-	1
7	+1	-	-	1	-	1	+1
8	+1	-	1	+1	1	-	+1
Allocation	3	5	1	2	1	1	3
Payments	20	9	8	15	8	7	20
Profits	3	26	2	13	7	5	33

Table 3.5: Clinching process in Example 3.3.1

mechanism customers have the incentive to bid truthfully in the lower-level auction, as was the case with the ACC auction in the non-hierarchical case. This applies despite the fact that in our case, the providers do not have a pre-specified supply, but gain incrementally in each occasion units of bandwidth that are simultaneously made available to their customers. Since no bid information is made available during the procedure, this increasing local supply is not known. Customer utility functions and thus local demand is independent of forthcoming increases of supply, thus implying that customers' strategies are not affected by the amount of bandwidth that is made available to them in each price. The clinching and payment rules for the lower level that we defined in Section 3.3, guarantees that each unit's standing price in the top level never exceeds that in the lower level. This guarantees non-negative profits for the providers (actually on a per unit basis) but raises further concerns about players' incentives. These are analyzed below.

3.4.1 Players' Dominant Strategies

Providers wish to maximize their profits from participation in two trading markets that interact with each other only through their own actions. Their strategy involves a buy process in the top level and a sell process in the lower level. Since it is taken that the social planner has pre-specified the mechanism in both levels including charging in the lower level, the provider's strategy reduces to the bidding strategy of the top level on the basis of the information progressively revealed to him in his own local market auction. The key factor determining his optimal bidding strategy in the top-level auction is his utility function. If the demand in his local market were known and he had the authority to define the payment rule for his own customers, he could calculate his utility function completely and bid truthfully according to this function in the top-level auction. In our setting neither local demand is known nor the payment rule is determined by the provider. Demand is derived step-by-step starting from lower prices (higher quantities) to higher prices (lower quantities). At each price the information available to the provider is the demand up to this level. On the other hand, each customer takes part in the local market auction and wishes to maximize his net benefit according to his known utility function.

We henceforth restrict attention to:

- the following strategy of providers, referred to as demand revelation: bid the quantity indicated by local demand if not exceeding the capacity C, otherwise bid C and

- the following strategy of customers, referred to as truthful bidding: bid the quantity indicated by own utility function at each price p if not exceeding the capacity C, otherwise bid C.

Recalling that $Q_j(p)$ and $x_i(p)$ denote the bidding strategies of providers and customers respectively, aforementioned strategies are given as follows:

$$Q_j(p) = \min\{C, \sum_{i \in S_j} x_i(p)\} \quad \text{and} \quad x_i(p) = \begin{cases} \max\{k : \theta_{i,k} > p\}, & \text{if } \theta_{i,1} > p \\ \\ 0, & \text{otherwise,} \end{cases} \tag{3.11}$$

where we have also used the fact that marginal values of the same customer are diminishing. We claim that provider j maximizes his profit by adopting $Q_j(p)$ and customer i maximizes his net benefit by adopting $x_i(p)$, regardless of the strategies of other players.

Proposition 3.4.1 *Demand revelation constitutes a weakly dominant strategy for every provider.*

Proof. Consider a fixed provider j. Suppose that each other provider bids according to an arbitrary but fixed strategy of his own and the same applies to each customer. We will prove that provider j maximizes his profit by revealing his local demand.

Suppose that at a given price, provider j bids a quantity *less* than this demand. The arguments that follow, hold for both the case of competitive market and the case of non-competitive market. This is due to the fact that the case of non-competitive market reduces to the case of competitive market if the provider who obtains the entire capacity bids a quantity that is less than this capacity. Since the other providers and provider's j customers do no deviate from their strategies, the procedure will remain the same except for termination, that will be reached earlier, i.e., at a lower price. Indeed, provider j will win less or the same number of bandwidth units than he would if he had revealed his demand. Due to the second restriction of the payment rule in the lower level, all units will be sold at most at this price. This implies that:

1. Provider's j buy price for all bandwidth units won is the same compared to the case where his bid equals his demand at each price $\rightarrow$ no extra profit.

2. His sell price for bandwidth units sold up to the new market-clearing price is the same $\rightarrow$ no extra profit from these units.

3. His sell price for bandwidth units that would be sold at higher prices is now the new market-clearing price, which is lower than the original one $\rightarrow$ loss.

4. He will not win some units of bandwidth that, a subset of which could be sold with positive profit $\rightarrow$ possible loss.

Conversely, suppose that at a certain price provider j bids a quantity *higher* than his demand, yet not exceeding the capacity C. This strategy is meaningful only for the case of competitive market. Termination of the procedure will be delayed, which implies that:

1. Provider j may possibly obtain more units that he will not be able to sell $\rightarrow$ possible loss.

2. His buy price for units besides the extra ones does not change since the other providers do not change their bids $\rightarrow$ no extra profit.

3. Besides the extra units, the other ones will not be sold in higher prices since his customers do not change their bidding $\rightarrow$ no extra profit.

This also applies in the case where provider j obtains the entire capacity due to overbidding (case of non-competitive market). The extra units won by the provider due to overbidding are are not demanded by his customers at all.

Thus, at each price, provider j should bid his local demand, regardless of the other players' strategies. ■

Proposition 3.4.2 *Truthful bidding constitutes a weakly dominant strategy for every customer.*

Proof. Suppose that each provider bids according to an arbitrary but fixed strategy of his own. We will prove that every customer maximizes his net benefit by bidding truthfully. Indeed, the number of a customer's clinched units of bandwidth at each price is *independent* of his own bids. The units of bandwidth and the prices at which he gains them depend on the bids of his competitors in the same market and on local supply, which is derived by the other providers' bids. In other words, if, at a certain price, customer i reports a higher quantity than the true one, then he might win an extra unit at a price that is higher than his corresponding marginal value, thus resulting in a loss. Conversely, if customer i reports a lower quantity than the true one, then he faces a loss from not winning his last unit with positive net benefit, without achieving a lower price for any of his other units won. Thus, customer i bids truthfully. ■

Corollary 3.4.1 *In the case of competitive market, if each provider's bidding strategy is demand revelation, then all auctions in which a positive portion of capacity is allocated, terminate simultaneously when demand equals supply in the top-level auction.*

Proof. At price p_{f} demand equals supply in the top-level auction. At the corresponding price each local market demand (i.e., the final bid of each provider) equals the respective local market supply (i.e., the amount he has gained, which in turn equals the provider's final bid in the top level), the value of which is only then determined. Thus, all lower-level auctions are terminated no later than price p_{f} as well.

Additionally, each "competitive" [2] provider sells at least one unit at price p_{f}: If at price p_{f} no customers of a specific provider clinches any units, this means that neither his customers

[2] A competitive provider is the provider who obtains a positive portion of the capacity.

changed bids at this price, nor did local supply change. That is, no customer of any "competitive provider" changed bids so the procedure should have already terminated. Clearly, this constitutes a contradiction. Thus, no lower-level auction (of a "competitive provider") has been terminated at a lower price than p_{f}.

Consequently, all lower-level auctions in which a positive portion of capacity is allocated, terminate simultaneously at price p_{f}. ■

Remark 3.1

The case at which one provider obtains the entire capacity and the demand in his local market equals the capacity at price p_{f}, is a special case of the non-competitive market. Obviously, corollary 3.4.1 holds for this special case.

3.4.2 Efficiency and Social Welfare

The selected mechanisms in both levels lead to efficient outcomes if considered independently. As we showed in Section 3.2.4, this does not suffice for achieving the overall efficiency in our hierarchical allocation problem. However, for the case of non-competitive market overall efficiency is achieved, since all capacity is allocated by the unique provider efficiently. Henceforth we restrict attention to the case of competitive market which is more interesting and harder to analyze. Revelation of local demand for every provider is a necessary condition for the mechanism to be efficient. This implies that provider's j marginal utility (revenues for the k^{th} additional unit) should equal the local market's marginal utility $v_{j,k}$. Our mechanism satisfies the condition so that the final price (which is the maximum price at which a won unit can be sold) be equal to each provider's marginal utility for the last demanded unit of bandwidth (provided that this provider has won at least one unit). Indeed, since the providers reveal demand, each of them will sell at least one unit at final price p_{f} for otherwise the procedure would have terminated earlier (see proof of Corollary 3.4.1). Thus, the final price equals each provider's revenues for the last unit of bandwidth, that is his marginal utility. Note that this is a less general condition than that imposed in Section 3.2. The provider need not reveal all the demand; a part of it starting from low prices up to the final one is sufficient to achieve efficiency.

Our mechanism has the important property that the social welfare attained in a market with intermediaries is the same as that under direct allocation by the social planner. More interesting is the fact that customers' net benefits are identical in both cases. However, the social planner faces a loss that is conveyed as profit to the providers. This is proved in the next proposition:

Proposition 3.4.3 *For each customer, allocation and payments arising under our hierarchical auction mechanism are the same with those when the trade is performed directly and efficiently by the social planner.*

Proof. Assume first that the social planner allocates C units of bandwidth directly to the customers. Let p_1 be the price at which customer 1 clinches his first unit of bandwidth. Then p_1 is the smallest value that satisfies the condition:

$$C - \sum_{i \neq 1} q_i(p_1) = 1 \Leftrightarrow C - \sum_{i=1}^{N} q_i(p_1) + q_1(p_1) = 1, \tag{3.12}$$

where $q_i(p_1)$ is the bid of customer i at price p_1, $i = 1, \ldots, N$.

We will prove that at the hierarchical auction we propose, customer 1 clinches his first unit of bandwidth at price p_1 too. Suppose there are two providers A and B, and customer 1 belongs to the set S_A of customers of provider A. Let $Q_A(p), Q_B(p)$ be the bids of providers A and B respectively and $a_A(p)$ the quantity already clinched by provider A at price p. Then

$$a_A(p) = C - Q_B(p). \tag{3.13}$$

Let $p_1^{'}$ be the price at which customer 1 now clinches his first unit of bandwidth. Then $p_1^{'}$ is the smallest value that satisfies the condition:

$$a_A(p_1^{'}) - \sum_{i \in S_A, i \neq 1} q_i(p_1^{'}) = 1 \Leftrightarrow a_A(p_1^{'}) - Q_A(p_1^{'}) + q_1(p_1^{'}) = 1.$$

Combining this with (3.13), we obtain

$$C - Q_B(p_1^{'}) - Q_A(p_1^{'}) + q_1(p_1^{'}) = 1 \Leftrightarrow C - \sum_{i=1}^{N} q_i(p_1^{'}) + q_1(p_1^{'}) = 1. \tag{3.14}$$

Combining this with (3.12) we obtain $p_1 = p_1^{'}$. Applying inductively the same argument it follows that for $k = 2, \ldots, x_1$, (where x_1 is the quantity customer 1 ultimately obtains), the price at which customer 1 clinches his k^{th} unit is the same in both mechanisms. ■

Corollary 3.4.2 *The hierarchical auction mechanism yields the efficient outcome.*

Proof. Allocation of bandwidth to customers under our hierarchical auction mechanism is the same with that under ACC, which is efficient. Hence, efficiency of our mechanism follows. ■

To conclude, we defined a mechanism to allocate bandwidth efficiently in the presence of intermediaries, who make profits too, taking advantage of the fact that demand is revealed in the right order due to the proposed coordination of the auctions in the two levels.

3.4.3 Implementation and Practical Issues

In the model presented in Section 3.3 we have assumed that auctions of both levels are performed simultaneously. To simplify implementation, we propose an alternative version of the mechanism: auctions of the lower level are performed asynchronously and prior to the top-level auction; each lower-level auction is rum until the total demand equals the total available bandwidth C. As soon as all providers have completed their respective own local auctions, the top-level auction is run. All the other rules and restrictions still apply.

Even though providers know their own local market's demand before entering the top-level auction, they are forced to apply the payment rule of Section 3.3 that induces truthful revelation of demand. Additionally, customers have the same set of information and the same payments in both implementations, implying that their strategies will be the same, that is truthful bidding. Since all players' strategies and allocation rules still apply under the asynchronous implementation, the outcome will be the same too.

This alternative implementation is less complicated because input from the various parts need not be gathered at the same time. However, it is required that each customer reveals his utility for the whole range of available bandwidth, rather than for the part of the bandwidth to be made available in his respective local market. This constitutes unnecessary revelation of information.

Another difficulty involved in our approach, is the implementation of the continuous clock that represents the price. Practically, only the prices at which there is a change in demand need to be considered. The set of these instances is finite since identical and indivisible units are being sold. Thus instead of reporting a continuous set of bids, it suffices to report only a finite set of bids.

Finally, it is important to guarantee that providers conform to the rules imposed by the social planner. The *auditability* of our approach is indeed attained by announcing the prices of all units sold in the top-level auction as well as in the lower-level auctions. Each customer can verify that his bids were treated as specified by the mechanism.

3.5 Competition

Thus far, we have assumed that the social planner regulates the bandwidth trade by imposing certain rules to the providers. In this section we will try to define the minimum set of rules and regulations to be imposed by the social planner in order to have his initial requirement for efficiency fulfilled. Obviously, if a provider is allowed to choose the mechanism for his own local market, then he will try to maximize his profits, without being concerned with social welfare. It is important to examine whether or not such a choice is in conflict with the proposed solution of Section 3.3, in which the social planner directs the whole system.

3.5.1 Model I: No Restrictions for Providers

Suppose that the social planner defines the top-level mechanism as described in Section 3.3 and imposes no rules to the providers on how each of them will charge his own market. Each market is assumed to have a known and fixed number of customers. Each provider seeks to maximize his profits by choosing an appropriate auction for the lower-level trade. However, he is restricted by the fact that he has to compete for bandwidth himself in the top-level auction based on the demand he will learn by the bids of his own customers.

The payment rule of the hierarchical auction mechanism we proposed in Section 3.3, is not optimal in terms of profits, for the provider, in the following sense: It is the restriction on the maximum permissible price in the lower level that renders demand revelation a weakly dominant strategy. Had it been omitted, providers would shade demand to their benefit: Lower-level auctions would terminate at higher prices (not simultaneously) yielding more revenues per unit won, while the top-level auction would terminate at a lower price yielding smaller charges, as shown in Example 3.5.1.

Provider A		**Provider B**	
Cust1	Cust2	Cust3	Cust4
$\theta_{\mathbf{1},\mathbf{1}} = \mathbf{12}$	$\theta_{2,1} = 7$	$\theta_{\mathbf{3},\mathbf{1}} = \mathbf{13}$	$\theta_{\mathbf{4},\mathbf{1}} = \mathbf{8}$
$\theta_{\mathbf{1},\mathbf{2}} = \mathbf{10}$	$\theta_{2,2} = 5$	$\theta_{\mathbf{3},\mathbf{2}} = \mathbf{11}$	$\theta_{4,2} = 4$
$\theta_{1,3} = 6$	$\theta_{2,3} = 3$	$\theta_{\mathbf{3},\mathbf{3}} = \mathbf{9}$	$\theta_{4,3} = 2$

Table 3.6: Marginal valuations in Example 3.5.1

	Lower Level A		**Lower Level B**		*Top Level*	
Price	Cust1	Cust2	Cust3	Cust4	ProvA	ProvB
1	3	3	3	3	6	6
2	3	3	3	2	6	5
3	3	2	3	2	5	5
4	3	2	3	1	5	4
5	3	1	3	1	4	4
6	2	1	3	1	3	4
7	2	-	3	1	2	4

Table 3.7: Bidding process in the original mechanism in Example 3.5.1

Example 3.5.1

Suppose there is an amount of $C = 6$ units of bandwidth that is allocated to two service providers A and B that have two customers each, with marginal valuations shown in Table 3.6.

In Tables 3.7 and 3.8 we present the bidding and clinching process of the original mechanism respectively.

We apply again the proposed mechanism relaxing the restriction on the maximum permissible price in the lower level. In Table 3.9 we present the bidding process. Each provider observes that his demand equals that of the other provider at price 1. Thus, it is plausible that both of them decide to reduce demand to 3 units, thus sharing evenly the bandwidth between them at a low price. We observe that provider's A local auction terminates at price 6 and provider's B local auction terminates at price 8, whereas top-level auction terminates at price 2. The allocation and

	Top level		**Lower level A**		**Lower Level B**	
Price	ProvA	ProvB	Cust1	Cust2	Cust3	Cust4
1	-	-	-	-	-	-
2	1	-	-	-	-	-
3	-	1	-	-	-	-
4	+1	-	-	-	-	-
5	-	+1	1	-	1	-
6	-	+1	-	-	+1	-
7	-	+1	+1	-	+1	1
Allocation	2	4	2	0	3	1
Payments	6	21	12	0	18	7
Profits	6	4	10	0	15	1

Table 3.8: Clinching process in the original mechanism in Example 3.5.1

	Lower Level A		**Lower Level B**		*Top Level*	
Price	Cust1	Cust2	Cust3	Cust4	ProvA	ProvB
1	3	3	3	3	6	6
2	3	3	3	2	3	3
3	3	2	3	2		
4	3	2	3	1		
5	3	1	3	1		
6	2	1	3	1		
7			3	1		
8			3	-		

Table 3.9: Bidding process in relaxed model in Example 3.5.1

	Top level		**Lower level A**		**Lower Level B**	
Price	ProvA	ProvB	Cust1	Cust2	Cust3	Cust4
1	-	-	-	-	-	-
2	3	3	-	-	1	-
3			1	-	-	-
4			-	-	+1	-
5			+1	-	-	-
6			-	1	-	-
7					-	-
8					+1	-
Allocation	3	3	2	1	3	0
Payments	6	6	8	6	14	0
Profits	8	8	14	1	19	0

Table 3.10: Clinching process in relaxed model in Example 3.5.1

payments of all players in this case are calculated in Table 3.10. We see that provider A improves his profits from 6 to 8 units of money and provider B improves his profits from 4 to 8 units of money. Thus, demand revelation is *not* a dominant strategy any more when the restriction on the maximum permissible price is relaxed. Moreover, this ruins efficiency of the allocation of bandwidth to customers. Recall, that the efficient allocation is the vector $(2, 0, 3, 1)$. ▲

We have shown that the hierarchical auction mechanism in not an optimal mechanism for providers in the unrestricted environment of model I. In fact, there can be found other auctions too in the lower level, that yield higher profits to the providers. Moreover, we have assessed the significance of the restriction of the maximum permissible price in the lower-level auctions. In the following, we will define models in which the hierarchical auction mechanism performs well in terms of providers' profits.

3.5.2 Model II: Restrictions to Providers' Highest Price

A somewhat more restrictive model for providers is now considered, yet less restrictive from the original model of Section 3.3. The social planner still controls the top-level auction which is again the Ascending Clock Auction with Clinching, and imposes the following price restriction to the providers for the lower-level auction:

> If the case of competitive market applies, the maximum permissible price for each unit of bandwidth to be sold by each provider is the final price of the top-level auction.

The lower-level auctions may take place either simultaneously with the top-level auction, or asynchronously prior to the top-level auction. The above price restriction is necessary, for if it is omitted the providers can find a more profitable strategy than truthfully revealing demand (see Example 3.5.1), which in general results in an inefficient outcome.

Under the new setting, the providers are asked to solve two different problems: First, a provider has to select an allocation mechanism in the lower level. After the whole mechanism is defined, the provider has to calculate the optimal (with respect to his profits) strategy for the top-level auction. These two problems are not independent and thus, they should be treated together. Each provider is now free to charge each unit of bandwidth at any price up to the maximum value imposed by the social planner. The pricing rule of the lower-level mechanism introduced in Section 3.3 and the uniform pricing rule (pay the market clearing price for all units obtained) both satisfy the above restriction. These pricing rules constitute the two *extreme* cases among all rules conforming to the above restriction: the former sets the lowest possible prices for each unit of bandwidth, since each such unit is made available in the lower-level auction as soon as it is won in the top-level auction; the latter sets the maximum price for all units of bandwidth. Next we compare these two schemes in terms of provider profits and examine under various assumptions whether our mechanism will indeed be employed by the providers. It is assumed that the payment rule is announced to the customers prior to trading.

Provider A		**Provider B**	
Cust1	Cust2	Cust3	Cust4
$\theta_{\mathbf{1},\mathbf{1}} = \mathbf{10}$	$\theta_{\mathbf{2},\mathbf{1}} = \mathbf{12}$	$\theta_{\mathbf{3},\mathbf{1}} = \mathbf{11}$	$\theta_{\mathbf{4},\mathbf{1}} = \mathbf{13}$
$\theta_{1,2} = 4$	$\theta_{\mathbf{2},\mathbf{2}} = \mathbf{5}$	$\theta_{\mathbf{3},\mathbf{2}} = \mathbf{9}$	$\theta_{\mathbf{4},\mathbf{2}} = \mathbf{8}$
$\theta_{1,3} = 1$	$\theta_{2,3} = 3$	$\theta_{3,3} = 2$	$\theta_{\mathbf{4},\mathbf{3}} = \mathbf{6}$

Table 3.11: Marginal valuations in Example 3.5.2

3.5.3 Comparison of Uniform Pricing and the Original Payment Rule with respect to Profits in the Lower Level

In this section we will compare two alternative payment rules applied to the lower level from the perspective of the provider profits. Assume that the providers charge their customers by applying *uniform pricing*: each unit is sold at the market clearing price. Providers' dominant strategy is to reveal the demand expressed through customers' bid, since the reasoning of Section 3.4.1 still holds. On the other hand, knowing the payment rule, customers have the incentive of demand reduction. In general, the outcome of this game will not be efficient and providers' profits will be in some cases higher and in other cases lower compared to our mechanism. Below we present Example 3.5.2 where uniform pricing is not beneficial for any provider, and Example 3.5.3 where uniform pricing is indeed beneficial for one of them. In both examples we assume complete availability of information. That is, customers' utility functions are common knowledge to all players.

Example 3.5.2

Suppose there is an amount of $C = 8$ units of bandwidth that is allocated to two service providers A and B that have two customers each, with marginal valuations shown in Table 3.11. In Tables 3.12 and 3.13 we present the bidding and clinching process of the original mechanism respectively.

Next, we apply uniform pricing in the lower level. Providers reveal demand truthfully. Moreover, it is proved in Appendix B that truthful bidding of all customers at each price except for customer 2 who reduces demand by 2 units at price 2, constitutes a Nash equilibrium in the complete information game. Note that other customers too, have the incentive to reduce demand due to the uniform pricing rule taking into consideration their opponents' bidding, but at different

	Lower Level A		**Lower Level B**		*Top Level*	
Price	Cust1	Cust2	Cust3	Cust4	ProvA	ProvB
1	2	3	3	3	5	6
2	2	3	2	3	5	5
3	2	2	2	3	4	5
4	1	2	2	3	3	5

Table 3.12: Bidding process in the original mechanism in Example 3.5.2

	Top level		**Lower level A**		**Lower Level B**	
Price	ProvA	ProvB	Cust1	Cust2	Cust3	Cust4
1	2	3	-	-	-	-
2	+1	-	-	1	-	1
3	-	+1	1	-	1	+1
4	-	+1	-	+1	+1	+1
Allocation	3	5	1	2	2	3
Payments	4	10	3	6	7	9
Profits	5	6	7	11	13	18

Table 3.13: Clinching process in the original mechanism in Example 3.5.2

	Lower Level A		**Lower Level B**		*Top Level*	
Price	Cust1	Cust2	Cust3	Cust4	ProvA	ProvB
1	2	3	3	3	5	6
2	2	**1**	2	3	3	5
Payments	4	2	4	6	4	7
Profits	10	10	16	21	2	3

Table 3.14: Uniform pricing in lower level in Example 3.5.2

	Top level	
Price	ProvA	ProvB
1	2	3
2	+1	+2
Allocation	3	5

Table 3.15: Clinching process in the top level under uniform pricing in the lower level in Example 3.5.2

prices each. The first customer to proceed to such a reduction is customer 2 at the price 2, which thus becomes the new final price of all auctions. The bidding process, the payments and profits are shown in Table 3.14. Note that providers' payments are calculated according to the clinching process of the top level, which is shown in Table 3.15. Both providers A and B have lower profits under the uniform pricing rule than under our mechanism. Moreover, the final allocation under the uniform pricing rule, is not the efficient one. ▲

Example 3.5.3

Assume that two providers A and B compete for $C = 5$ units of bandwidth. Each provider has two customers. Their marginal utilities are shown in Table 3.16. The efficient outcome is given by vector $\vec{x}^* = (x_1^*, \ldots, x_4^*) = (4, 0, 1, 0)$. We apply the hierarchical auction using the uniform pricing rule in the lower level. Reasoning as in Appendix B, it can be checked that the equilibrium strategy of customers is as follows: All customers bid truthfully at each price except for customer 1

Provider A		**Provider B**	
Cust1	Cust2	Cust3	Cust4
$\boldsymbol{\theta_{1,1} = 16}$	$\theta_{2,1} = 8$	$\boldsymbol{\theta_{3,1} = 11}$	$\theta_{4,1} = 9$
$\boldsymbol{\theta_{1,2} = 14}$	$\theta_{2,2} = 8$	$\theta_{3,2} = 7$	$\theta_{4,2} = 4$
$\boldsymbol{\theta_{1,3} = 12}$	$\theta_{2,3} = 6$	$\theta_{3,3} = 5$	$\theta_{4,3} = 4$
$\boldsymbol{\theta_{1,4} = 10}$	$\theta_{2,4} = 6$	$\theta_{3,4} = 5$	$\theta_{4,4} = 4$

Table 3.16: Marginal valuations in Example 3.5.3

	Lower Level A		**Lower Level B**		*Top Level*	
Price	Cust1	Cust2	Cust3	Cust4	ProvA	ProvB
4	4	4	4	1	5	5
5	4	4	2	1	5	3
6	4	2	2	1	5	3
7	4	2	1	1	5	2
8	**3**	-	1	1	3	2
Allocation	3	0	1	1	3	2
Payments	24	0	8	8	17	16
Profits	18	0	3	1	7	0

Table 3.17: Uniform pricing in lower level in Example 3.5.3

who reduces demand by 1 unit at price 8. Providers again reveal demand truthfully. The bidding process, the payments and profits are shown in Table 3.17. Note that providers' payments are calculated according to the clinching process of the top level: Provider A clinches two units at price 5 and one unit at price 7; he thus pays an amount of $2 * 5 + 7 = 17$ units. Provider B clinches two units at price 8, paying an amount of $2 * 8 = 16$ units. Players' bidding process and profits had we applied the original mechanism are given in Tables 3.18 and 3.19 respectively. We observe that provider A (whose customer reduced demand) has higher profits (7 versus 6) under the uniform pricing rule. Moreover, the outcome is not efficient. ▲

	Lower Level A		**Lower Level B**		*Top Level*	
Price	Cust1	Cust2	Cust3	Cust4	ProvA	ProvB
4	4	4	4	1	5	5
5	4	4	2	1	5	3
6	4	2	2	1	5	3
7	4	2	1	1	5	2
8	4	-	1	1	4	2
9	4	-	1	-	4	1

Table 3.18: Bidding process in the original mechanism in Example 3.5.3

	Top level		**Lower level A**		**Lower Level B**	
Price	ProvA	ProvB	Cust1	Cust2	Cust3	Cust4
4	-	-	-	-	-	-
5	2	-	-	-	-	-
6	-	-	-	-	-	-
7	+1	-	1	-	-	-
8	-	1	+2	-	-	-
9	+1	-	+1	-	1	-
Allocation	4	1	4	0	1	0
Payments	26	8	32	0	9	0
Profits	6	1	20	0	2	0

Table 3.19: Clinching process in the original mechanism in Example 3.5.3

These examples suggest that the model of Section 3.5.2 where providers are less restricted than in the initial model of Section 3.2, does not in general lead to the efficient solution. Providers take advantage of their chance to select the most profitable way to charge their customers and, as we saw, there are cases where the outcome differs from the efficient one.

3.5.4 Model III: Relaxing Restrictions for Customers

Thus far, we have assumed that customers are committed to a local market. We extend the model of Section 3.5.2 by letting customers *choose* a provider after they become aware of the payment rules of all providers. As in our initial model, after local markets are formed, customers are not allowed to move to another market. Obviously, customers take advantage of the information on the payment rules employed and choose the provider from which they will obtain the highest net benefit. Providers' strategy now becomes more complicated since they should take into account customers' freedom of market selection. On one hand, uniform pricing might lead a provider to higher profits (yet not always) had his population of customers not changed. On the other hand, as a reaction to choosing uniform pricing, the population of this provider's customers might shrink, thus leading him to lower profits. Actually *no* customer will select this provider as we argue below! Next, we establish some very interesting properties that support the selection of the original mechanism by the providers.

Proposition 3.5.1 *Under model III, if all providers choose the payment rule for the lower level that is defined in Section 3.3, then each customer has the same net benefit in any market he chooses and for any distribution of the rest of the customers among the markets.*

Proof. Since the bidders' marginal valuations for each unit of bandwidth are assumed to be drawn from a continuous set of values, the probability of two or more bidders having the same valuation for a unit is negligible. This implies that the efficient allocation is unique. Consequently, all possible placements of customers in the various markets result in the same outcome, namely, the unique efficient allocation. The result follows by combining this with Proposition 3.4.3. ■

Proposition 3.5.1 implies that customers are indifferent regarding which provider to choose, assuming that all providers apply the same charging scheme. In the sequel, we examine customers' choices in case a provider, say A, has announced the selection of uniform pricing as the payment rule in his local market's auction and the other providers have selected the original mechanism.

Proposition 3.5.2 *Under model III, any fixed customer i has the incentive not to choose the provider applying uniform pricing.*

Proof. Assume that provider's A market consists of group S_A of customers and all other providers have the group S' of customers. Consider the fixed customer i. We remind that truthful bidding is customers' weakly dominant strategy under our mechanism. Thus, customers of all providers apart from A will still bid truthfully. We define three possible scenarios:

1. Customer i chooses group S_A, i.e. joins the local market of provider A. Due to uniform pricing, his optimal strategy is to shade bids. Let $NB_i^{(1)}$ be his net benefit.

2. Customer i chooses group S' but employs the same strategy as in scenario 1. We will show that his new net benefit $NB_i^{(2)}$ is at least as much as previously, i.e. $NB_i^{(1)} \leq NB_i^{(2)}$. First note that all customers not belonging to S_A bid truthfully as in scenario 1, due to Proposition 3.4.2. Providers bid truthfully in both scenarios as well, since they do not depend on their respective populations of customers. Consequently, total demand remains the same in each price in each level, so the final price of the top level will be the same in both scenarios. Again, since all players apply the same strategies, the final allocation will be the same. Thus, customer i wins the same units of bandwidth, each at a price less than or equal to the final price of scenario 1 (which is equal to the final price of scenario 2). This results from the fact that the providers of group S' charge according to the clinching prices of the original mechanism that are less than the final price.

   Consider the special case where S_A consists only of customer i who subsequently leaves this market, and the group S' belongs to one provider, say B. Then, the non-cooperative case arises. In this case, $NB_i^{(1)} \leq NB_i^{(2)}$ still holds: if customer i wins no bandwidth in group S_A, then he wins no bandwidth in group S'; thus, the equality holds: $NB_i^{(1)} = NB_i^{(2)}$. If customer i obtains a positive portion of capacity from provider A, then he obtains this quantity from provider B too. This implies that provider's B demand (not just his bid) equals supply at price p_{f}. Thus, the results of the competitive market still apply according to the remark in Section 3.4.1. Therefore, $NB_i^{(1)} \leq NB_i^{(2)}$ in any case.

3. Customer i chooses group S' and plays according to his (weakly) dominant strategy; that is, he bids truthfully. Consequently, his net benefit $NB_i^{(3)}$ is greater than or equal $NB_i^{(2)}$.

This together with the previous conclusion implies that $NB_i^{(1)} \leq NB_i^{(3)}$. Therefore, customer i will choose group S', i.e. he will not join the local market of provider A. ■

Thus, no customer will choose provider A who announced to have selected uniform pricing. Combining this result with the property given in Proposition 3.5.1, each customer will arbitrarily choose one of the providers adopting the original mechanism. Consequently, no provider has the incentive to choose uniform pricing. In other words, under model III, each provider is better off in terms of profits with the original mechanism.

3.6 Concluding Remarks

In this chapter, we focus on strategic interactions among sellers, retailers and potential buyers. We propose and analyze a new hierarchical auction to allocate bandwidth efficiently in a hierarchical structured market. We take advantage of the distribution of information over the parts involved and coordinate the various trades that have to take place, such that no one has the incentive to deviate from bidding truthfully. We prove that despite the consecutive transactions, customers do not incur further losses (this is the cost the social planner has to pay to assure efficiency). This is an important property, because the hierarchical structure does not affect customers. We also argue that applying efficient mechanisms in each level, this alone does not guarantee overall efficiency. The key issue is demand revelation by the providers who are enforced to do so by restricting their choices.

We also define business models that differ in the distribution of information and the level of decision making each player possesses, and we investigate the impact competition has on their profits. The above models are based on the initial hierarchical auction used to sell bandwidth in multiple levels efficiently. We show that the more power an intermediate provider has, the more profits can he extract from his customers. However, we also prove that if each customer is allowed to choose his own provider on the basis of the selected payment rules, then each provider has the incentive to apply the initial payment rule, for otherwise he would end up with no customers.

Our results are not particular for bandwidth markets; they also apply to other markets for which the trading hierarchical model pertains, such the hierarchical trading of units of other services (e.g. call minutes), or the trading of bandwidth of an overprovisioned backbone link (top-level trade) that interconnects with congested access networks (lower-level trade), considered by Maillé and Tuffin in [24].

Chapter 4

Seller Participation in First and Second Price Auctions

4.1 Introduction

An issue that has received attention in auction theory and practice is the influence that the seller of a good may exert on the outcome of an auction, simply by choosing the auction type that would serve his own purposes. Even in regulated environments, where the type of auction to be employed is given by a third party, the seller may play an active role in the auction's progress. In many situations, he participates in an auction by setting a reserve price publicly announced to the bidders prior to the auction, or he retains the right to bid as if he were a bidder by himself. The latter perspective can be thought of as that of imposing a secret reserve price. We assume that bidders are informed of the seller's participation as opposed to the problem of "phantom" bidding in which the seller submits bids unofficially in order to raise the price. Bidding behavior is more complicated in this setting, since bidders have to take into account the behavior of the seller too.

The seller has the incentive to announce a reserve price in an effort to raise his expected profit. Indeed, it is proved in [25] that, in the case of symmetric bidders, the optimal [1] auction for a single good is any auction in which the seller keeps the good if the highest bid is less than the optimal reserve price; otherwise the bidder with the highest bid wins the good. In this context,

[1]Recall that an optimal auction is the auction that yields the highest expected profit for the seller.

first and second-price auctions with the corresponding optimal reserve price both maximize seller's expected profit. Moreover, the seller has the incentive to participate in an auction for a good that is valuable to him too. This may be the case when the seller is a service provider. If the bidders are not willing to pay an amount more than the seller's valuation, then the seller benefits by keeping it for personal use.

In this chapter we study the case of imposing a secret reserve price and whether this affects first and second-price auctions in terms of seller's expected profit, under the private values model. We show that in some cases a first-price auction is more profitable while in other cases a second-price auction yields higher profits to the seller. We also show that this does not contradict the fact that the second-price auction with reserve price is the optimal mechanism with respect to seller profits. In our model the first-price auction could yield more profits to the seller than the second-price in specific cases, provided that he would be forced to apply a first-price auction, so that the distribution of his valuation would be known to the bidders. If the seller has a choice of mechanism then the players' strategies that result in the aforementioned comparison of profits, do not constitute an equilibrium any more. This applies because the distribution of the seller's valuation, conditional on his selection of mechanism, affects bidders' strategies and should be derived as part of the equilibrium.

Chakraborty and Kosmopoulou deal with the problem of seller participation under the common values model in [5]. The authors of [5] analyze an ascending price bidding process for a single good and examine the effect seller participation has on profits. They prove that the potential of seller participation makes the seller worse off. In fact, any out-of-auction mechanism that makes it difficult for the seller to submit a bid, increases his revenues. In addition, they prove that the bidders' surplus and the surplus from trade is reduced.

Wolfstetter considers in [35] first price auctions with the seller keeping his reserve price secret. The author of [35] conjectures that bidders' equilibrium strategy is independent of the secret reserve price, and proves that the seller's equilibrium reserve-price strategy is to set it equal to the seller's own valuation. In our analysis we prove that it is a dominant strategy for the seller to set the secret reserve price equal to his valuation and that bidders' strategy is actually affected by the seller's presence.

Before proceeding with our analysis, we summarize the results of first and second-price auctions with a publicly known reserve price below. For a thorough study see [16].

4.1.1 Reserve Prices

Consider the problem of selling a single good to a set of N bidders. We examine the first and the second-price auction as trading mechanisms. The seller imposes a reserve price r, so that the good is not sold if the highest bid is less than r. The reserve price r is announced to all bidders at the beginning of the auction. Each bidder i has a privately known valuation x_i for the good. The seller has a valuation x_s which is his maximum willingness-to-pay for the good. That is, he wishes to keep the good for *himself* if no bidder is willing to pay more than x_s. Each bidder i is assumed to know the distribution functions of his rival's valuations $X_1, \ldots, X_{i-1}, X_{i+1}, \ldots, X_N$ respectively as well as the seller's valuation X_s. We assume that $X_1, \ldots, X_N$ and X_s are drawn independently and identically from the uniform distribution on $[0, 1]$: the cumulative distribution function equals $F(x) = x,\ x \in [0, 1]$ and the density function equals $f(x) = 1,\ x \in [0, 1]$. The seller knows the distribution functions of bidders' valuation too. We examine bidding behavior and how seller's expected profit is affected.

First-price auction with reserve price. Recall that in a first-price (sealed-bid) auction with reserve price, the winner is the bidder with the highest bid provided that this bid is higher than the reserve price r. The winner pays his own bid. All bidders have the same information for the game. Thus, there exists a symmetric equilibrium strategy β for each bidder. $\beta(x)$ denotes the bid submitted by a bidder whose valuation is x. If a bidder has a valuation $x < r$, then there is no possibility of winning the good. If a bidder has a valuation $x = r$, then $\beta(r) = r$. Indeed, if $\beta(r) < r$, then the bidder wins with probability 0. If $\beta(r) > r$, then he obtains a negative net benefit in case of winning. For all $x > r$, there holds $\beta(x) \geq r$.

It is proved in [16] that the symmetric equilibrium strategy for any bidder with valuation $x \geq r$ is as follows:

$$\beta(x) = \frac{(N-1) \cdot x^N + r^N}{N \cdot x^{N-1}}. \tag{4.1}$$

The optimal reserve price r^* and the seller's expected revenue E_s^* are:

$$r^* = \frac{x_s + 1}{2}, \text{ and} \tag{4.2}$$

$$E_s^* = \frac{(x_s + 1)^{N+1}}{2^N \cdot (N+1)} + \frac{N-1}{N+1}. \tag{4.3}$$

Note that the optimal reserve price is higher than the seller's valuation. According to [25], the first-price auction with reserve price is an optimal auction in the symmetric case. The underlining intuition is that the announcement of a reserve price induces more aggressive bidding so that for each fixed bidder the probability of winning be increased. Indeed, the bidders' strategy of Equation (4.1) gives a higher bid than that of a standard first-price auction given in Example 2.4.1. The seller benefits from the reserve price, since higher bids produce higher revenues too. However, the bidding strategy given in Equation (4.1) does *not* preserve the order of the bidders' valuations, resulting in an inefficient outcome. For example, let $N = 3$, $x_1 = 0.8$, $x_2 = 0.5$, $x_3 = 0.6$ and $x_s = 0.4$. Then $r = 0.7$ and the bidders submit the following bids: $b_1 = 0.7119$, $b_2 = 0.79$, $b_3 = 0.7175$. The winner is bidder 2 which is the bidder with the lowest valuation. Therefore, the first-price auction with reserve price is not efficient.

Second-price auction with reserve price. In a second-price (sealed-bid) auction truthful bidding is a weakly dominant strategy: $\beta(x) = x$ when a reserve price is introduced. The winner is the bidder with the highest bid provided that this bid is higher than the reserve price r. The winner pays the second highest bid provided that it is higher than the reserve price. Otherwise, he pays the reserve price. Thus, bidding behavior is not affected in the presence of a reserve price. Note that if $x < r$ the respective bidder will definitely not be the winner. It is proved in [16] that both the seller's expected profit and the optimal reserve price are the same as in the first-price auction for any distribution function of the bidders' valuation. Therefore, the second-price auction with reserve price and symmetric bidders is an optimal auction too. The fact that the reserve price might be higher than the second highest valuation results in a higher expected profit for the seller. Nevertheless, if no bidder has a valuation higher than the reserve price, the good is not sold even though there might be a bidder willing to pay more than the seller's valuation. Thus, second-price auction with reserve price is not efficient either. For example, let $N = 3$, $x_1 = 0.85$, $x_2 = 0.7$, $x_3 = 0.5$ and $x_s = 0.8$. Then, since bidders are truthful and $r = 0.9$, the good is not sold. Even though bidder 1 is the highest bidder, he does not obtain the good due to the high reserve price.

In conclusion, profit equivalence does hold in first and second-price auctions in the case of reserve prices for any distribution function of the bidders' valuation. Moreover these mechanisms are optimal for the seller in the case of symmetric and uniformly distributed bidders. In the next section, we study first and second-price auctions with seller participation and show that profit equivalence is no more valid.

4.2 Seller Participates in the Auction as a Bidder

We extend the model of Section 4.1.1 to the case where the seller competes for the good as if he were a bidder like the others. A bidder's strategy consists of the bid that denotes how much he is willing to "pay" for the good. The good is allocated to the bidder submitting the highest bid if that bid exceeds the one of the seller too. The seller's strategy consists of two parts: a) the type of the auction to be employed and b) the bid he will submit in the auction selected. In order to derive his optimal bid, a bidder now takes into consideration: 1) the auction rules, 2) his belief about the other bidders' and the seller's valuations, 3) the other bidders' strategies and 4) the seller's strategy. This last feature complicates bidding behavior substantially. The seller's choice of an auction type may reveal information about his valuation. Bidders anticipate the behavior of the seller adjusting their strategies.

Recall that we have assumed the private values model and that the distribution functions of all players' valuations are independently and identically drawn from the uniform distribution on interval [0,1]. We further assume that the seller is enforced to employ a first or a second-price auction. Thus, his strategy only consists of the bid he will submit in the auction and is derived by maximizing his expected profit:

$$E_s = P_s \cdot x_s + (1 - P_s) \cdot y_p, \tag{4.4}$$

where P_s is the probability of the seller keeping the good, and y_p is the payment of the winner in case the good is sold to a bidder. In a first-price auction y_p equals the highest bid, whereas in a second-price auction y_p equals the second highest bid including the seller's bid too. We derive equilibrium strategies for both mechanisms below.

4.2.1 First-Price Auction

We consider that the seller is enforced to employ the first-price auction, i.e. there is no other option of a mechanism. This means that the seller's strategy reduces to the choice of the bid he intends to submit. Thus, bidders obtain no extra information about the seller's valuation. They hold their initial belief that the seller's distribution function of his valuation is the uniform distribution on interval [0,1]. Next, we prove that:

1. *truthful bidding* is a dominant strategy for the seller, provided that bidders adopt a symmetric strategy.

2. bid shading, the same as in a classical first-price auction with $N+1$ bidders (i.e., the N original ones and the seller), is an equilibrium strategy for every bidder.

Wolfstetter considers in [35] this game and only proves the weaker result that it is an equilibrium strategy for the seller to bid truthfully. In order to prove this, he considers that bidders' strategy is not affected by the seller's presence, which, according to our analysis does not apply. To the best of our knowledge the analysis and the results to follow are new.

Proposition 4.2.1 *In a first-price auction with seller participation, truthful bidding is a dominant strategy for the seller provided that bidders adopt a symmetric strategy.*

Proof. Let the seller with valuation x_s submit a bid b_s. Assume further that each bidder with valuation x applies the symmetric, increasing and differentiable strategy $\beta(x)$. Each bidder's valuation is uniformly distributed on [0,1]. We will prove that the seller's optimal strategy β_s^f is to bid truthfully, against any β. We have

$$\Pr[\beta(X_j) < b_s] = \begin{cases} \beta^{-1}(b_s), & \text{for } b_s \leq \beta(1) \\ 1, & \text{for } b_s > \beta(1) \end{cases} \tag{4.5}$$

The seller's probability of winning is a function of his bid b_s and is given by:

$$\begin{aligned} P_s(b_s) &= \Pr[\max_{j=1,\dots,N} \{\beta(X_j)\} < b_s] \\ &= \Pr[\beta(X_1) < b_s \text{ and } \cdots \text{ and } \beta(X_N) < b_s] \\ &= \prod_{j=1}^{N} \Pr[\beta(X_j) < b_s] \\ &= \prod_{j=1}^{N} \Pr[X_j < \beta^{-1}(b_s)] \Leftrightarrow \\ P_s(b_s) &= \begin{cases} [\beta^{-1}(b_s)]^N, & \text{for } b_s \le \beta(1) \\ 1, & \text{for } b_s > \beta(1) \end{cases} \end{aligned} \tag{4.6}$$

where we have also used independence of $X_1, \dots, X_N$. The seller's expected profit E_s equals his valuation x_s if he is the winner; otherwise, it equals the expected value of the maximum bid Y of the N original bids conditional on the seller not being the winner. Therefore,

$$E_s(b_s) = \begin{cases} P_s(b_s) \cdot x_s + [1 - P_s(b_s)] \cdot E[Y|Y > b_s], & \text{for } b_s \le \beta(1) \\ x_s, & \text{for } b_s > \beta(1) \end{cases} \tag{4.7}$$

where $Y = \max_{j=1,\dots,N} \{\beta(X_j)\}$. Henceforth, to simplify notation, we use $\max_j \{\cdot\}$ instead of $\max_{j=1,\dots,N} \{\cdot\}$. Assume for now that $b_s \le \beta(1)$. The distribution function of Y given that $Y > b_s$, is

$$\Pr[\max_j \{\beta(X_j)\} \le y | \max_j \{\beta(X_j)\} > b_s] = \frac{\Pr[\beta^{-1}(b_s) \le \max_j \{X_j\} \le \beta^{-1}(y)]}{\Pr[\max_j \{X_j\} > \beta^{-1}(b_s)]} \tag{4.8}$$

Reasoning as in the case of (4.6), it follows easily that

$$\Pr[\max_j \{\beta(X_j)\} \le y | \max_j \{\beta(X_j)\} > b_s] = \frac{[\beta^{-1}(y)]^N - [\beta^{-1}(b_s)]^N}{1 - [\beta^{-1}(b_s)]^N}, \text{ for any } y \in [b_s, \beta(1)]. \tag{4.9}$$

Differentiating (4.9) with respect to y, we obtain the probability density function f_Y of Y conditional on $Y > b_s$. Thus, for the expected value at large, we have

$$\begin{aligned} E[Y|Y > b_s] &= \int_{b_s}^{\beta(1)} y \cdot f_Y(y)\, dy = \int_{b_s}^{\beta(1)} y \cdot \frac{N[\beta^{-1}(y)]^{N-1}}{1 - [\beta^{-1}(b_s)]^N} \cdot \frac{1}{\beta'(\beta^{-1}(y))}\, dy \\ &= \frac{N}{1 - [\beta^{-1}(b_s)]^N} \int_{b_s}^{\beta(1)} y \cdot [\beta^{-1}(y)]^{N-1} \cdot \frac{1}{\beta'(\beta^{-1}(y))}\, dy. \end{aligned} \tag{4.10}$$

Combining (4.6) and (4.7) with (4.10), the seller's expected profit is given by:

$$E_s(b_s) = [\beta^{-1}(b_s)]^N \cdot x_s + \{1 - [\beta^{-1}(b_s)]^N\} \cdot \frac{N}{1 - [\beta^{-1}(b_s)]^N} \int_{b_s}^{\beta(1)} y \cdot \frac{[\beta^{-1}(y)]^{N-1}}{\beta'(\beta^{-1}(y))}\, dy$$

$$= [\beta^{-1}(b_s)]^N \cdot x_s + N \int_{b_s}^{\beta(1)} y \cdot [\beta^{-1}(y)]^{N-1} \cdot \frac{1}{\beta'(\beta^{-1}(y))}\, dy. \tag{4.11}$$

In order to maximize E_s with respect to b_s, we differentiate the above expression. We have

$$E'_s(b_s) = N\frac{[\beta^{-1}(b_s)]^{N-1} x_s}{\beta'(\beta^{-1}(b_s))} - N\frac{b_s[\beta^{-1}(b_s)]^{N-1}}{\beta'(\beta^{-1}(b_s))}$$

$$= (x_s - b_s)N\frac{[\beta^{-1}(b_s)]^{N-1}}{\beta'(\beta^{-1}(b_s))}. \tag{4.12}$$

Therefore,

$$E'_s(b_s) = 0 \Leftrightarrow x_s = b_s, \tag{4.13}$$

while this point corresponds to a maximum because all functions β^{-1} and their derivatives are positive. If indeed $x_s \leq \beta(1)$, then $E_s(b_s)$ as given by (4.7) is maximized for $x_s = b_s$. Notice that the topmost expression in the right-hand side of (4.7) equals x_s for $b_s = \beta(1)$, which implies that if $x_s \leq \beta(1)$, then bidding a quantity $b_s > \beta(1)$ is not beneficial due to (4.13). On the other hand, if $x_s > \beta(1)$, then $E(b_s)$ is increasing in $[0, \beta(1)]$ and constant for $b_s > \beta(1)$. Therefore, any bid $b_s \geq \beta(1)$ is optimal for the seller which includes the case $b_s = x_s$. Therefore, in any case $\beta_s^f(x) = x$. That is, truthful bidding of the seller maximizes his expected profit. ■

Remark 4.1

Consider the more general case were each bidder i adopts a non-symmetric strategy β_i in a first-price auction with seller participation. Then, truthful bidding is *still* a dominant strategy for the seller provided that $\beta_1(1) = \cdots = \beta_N(1)$. The proof of this result is provided in Appendix C. In fact, we conjecture that truthful bidding by the seller applies for every strategy of the bidders. The proof of the general result should be similar to the Appendix C.

Proposition 4.2.2 *In a first-price auction with seller participation, the bidders' symmetric equilibrium strategy is* $\beta^f(x) = \frac{N}{N+1} \cdot x$.

Proof. Suppose that each bidder applies the symmetric, increasing and differentiable strategy $\beta(x)$. Consider a fixed bidder i with valuation x who submits a bid that equals b. From Proposition 4.2.1, bidder i knows that the seller applies the strategy $b_s = \beta_s^f(x) = x$. Moreover, bidder i assumes that the seller's valuation is uniformly distributed on [0,1]. We will prove that bidder's i optimal strategy is $\beta(x) = \beta^f(x) = \frac{N}{N+1} \cdot x$.

Bidder's i probability of winning equals

$$\begin{aligned} P_i(b) &= \Pr[\max_{j \neq i} \{\beta(X_j)\} < b \text{ and } \beta_s^f(X_s) < b] \\ &= \Pr[\max_{j \neq i} \{\beta(X_j)\} < b] \cdot \Pr[\beta_s^f(X_s) < b] \\ &= \left\{ \prod_{j \neq i} \Pr[\beta(X_j) < b] \right\} \cdot \Pr[X_s < b] \\ &= \left\{ \prod_{j \neq i} \Pr[X_j < \beta^{-1}(b)] \right\} \cdot \Pr[X_s < b] \\ &= (\beta^{-1}(b))^{N-1} \cdot b. \end{aligned} \tag{4.14}$$

Bidder's i expected net benefit E_i equals

$$E_i(b) = P_i(b) \cdot (x - b) = (\beta^{-1}(b))^{N-1} \cdot b \cdot (x - b). \tag{4.15}$$

Since β is a symmetric equilibrium strategy, there holds $b = \beta(x) \Leftrightarrow \beta^{-1}(b) = x$. In order to maximize E_i with respect to b, we differentiate $E_i(b)$. It follows from (4.15) that

$$\begin{aligned} E_i'(b) = 0 &\Leftrightarrow (b \cdot (\beta^{-1}(b))^{N-1})'(x - b) - b(\beta^{-1}(b))^{N-1} = 0 \\ &\Leftrightarrow [(\beta^{-1}(b))^{N-1} + b \cdot (N-1) \cdot \frac{(\beta^{-1}(b))^{N-2}}{\beta'(\beta^{-1}(b))}] \cdot (x - b) = b(\beta^{-1}(b))^{N-1} \\ &\Leftrightarrow [\beta^{-1}(b) + b \cdot (N-1) \cdot \frac{1}{\beta'(\beta^{-1}(b))}] \cdot (x - b) = b \cdot \beta^{-1}(b) \\ &\Leftrightarrow [x + b \cdot (N-1) \cdot \frac{1}{\beta'(x)}] \cdot (x - b) = b \cdot x. \end{aligned} \tag{4.16}$$

Solving the above differential equation with the initial condition $\beta(0) = 0$, we have

$$\beta(x) = \frac{N}{N+1} \cdot x. \tag{4.17}$$

Thus, the equilibrium strategy for bidder i equals $\beta^f(x) = \frac{N}{N+1} \cdot x$

■

The bidders' strategy is indeed affected by the presence of another "bidder", namely the seller. In particular, they consider him as being one of them. The bidder's equilibrium strategy has the same form with that of a first-price auction. However, the number of bidders N is just replaced by $N+1$. On the other hand, the seller's strategy is different from that of the bidders. This reflects the fact that the seller's objective function (i.e., expected profit) is different from the bidders' objective function (i.e., net benefit).

Seller's expected profit in first-price auction. Recall that the seller's expected profit when his bid equals b_s is $E_s(b_s)$ given by Equation (4.11). We derived in the proof of Proposition 4.2.1, that E_s is maximized at $b_s = x_s$ for any symmetric strategy of the bidders. However, the bidders will bid according to their equilibrium strategy b^f and will affect the seller's profit through this strategy. In order to calculate the optimal value of E_s we need to replace in Equation (4.11), b_s with x_s as well as the bidders' strategy β with the equilibrium strategy β^f. The optimal expected profit in a first-price auction is then given by:

$$\begin{aligned} E^f_{opt}(x_s) &= \left(\frac{N+1}{N}\right)^N \cdot x_s^N \cdot x_s + N\int_{x_s}^{\frac{N}{N+1}} y \cdot \left(\frac{N+1}{N}\right)^{N-1} \cdot y^{N-1} \cdot \frac{(N+1)}{N}\,dy \\ &= \left(\frac{N+1}{N}\right)^N x_s^{N+1} + \left(\frac{N}{N+1}\right)^2 - \left(\frac{N+1}{N}\right)^{N-1} x_s^{N+1} \qquad (4.18) \\ &= \frac{1}{N}\left(\frac{N+1}{N}\right)^{N-1} x_s^{N+1} + \left(\frac{N}{N+1}\right)^2 . \end{aligned}$$

4.2.2 Second-Price Auction

Next, we assume that the seller is enforced to employ a second-price auction, in which he participates too. Thus, his strategy consists of the bid he submits in the auction. Again, bidders' belief for the seller's distribution function is that it is drawn from the uniform distribution on [0,1].

In a second-price auction, bidders's weakly dominant strategy is truthful bidding. When the seller participates in a second-price auction, truthful bidding is *still* bidders' weakly dominant strategy: one more player (the seller) does not affect bidders' strategies no matter what his bid or valuation may be. Next, we prove that overbidding the same amount as in a second-price auction with reserve price is an equilibrium strategy for the seller.

Proposition 4.2.3 *In a second-price auction with seller participation, the seller's equilibrium strategy equals* $\beta_s^s(x) = \frac{x_s+1}{2}$.

Proof. Assume that the seller's valuation is x_s and his bid is b_s. He knows that each bidder bids according to the strategy $\beta(x) = x$ and that each bidder's valuation is uniformly distributed on [0,1]. The seller's probability of winning is

$$P_s(b_s) = \Pr[\max_{j=1,\ldots,N} \{\beta(X_j)\} < b_s] = [\beta^{-1}(b_s)]^N = b_s^N, \tag{4.19}$$

which follows reasoning similarly with (4.6). (We have assumed that $\beta_s^s(x) \leq 1$ for every $x \in [0,1]$, since bidding higher than 1 is equally good for the seller as bidding exactly 1.) If one of the N original bidders wins, then the payment depends on whether the seller has submitted the second highest bid. If this occurs, the winner pays the seller's bid, otherwise the winner pays the second highest bid among the other bidders' bids. Therefore the probability $P_s^{(2)}$ that the seller has submitted the second highest bid equals

$$\begin{aligned} P_s^{(2)}(b_s) &= \sum_{i=1,\ldots,N} \Pr[\beta(X_i) > b_s \text{ and } \max_{j \neq i} \{\beta(X_j)\} < b_s] \\ &= \sum_{i=1,\ldots,N} (1-b_s)b_s^{N-1} = N(1-b_s)b_s^{N-1}. \end{aligned} \tag{4.20}$$

The seller's profit equals: 1) x_s if he does not sell the good (i.e. if b_s is the highest bid), 2) b_s if he sells the good and his own bid is the second highest bid, 3) the second highest among the bidders' bid if b_s is less than the second highest bid. Thus, the seller's expected profit equals

$$E_s(b_s) = P_s(b_s) \cdot x_s + P_s^{(2)}(b_s) \cdot b_s + (1 - P_s(b_s) - P_s^{(2)}(b_s)) \cdot E[Y_2|Y_2 > b_s], \tag{4.21}$$

where Y_2 is the second maximum value among $\beta(X_1), \ldots, \beta(X_N)$. The distribution function of Y_2 equals $F_2(y) = Nb_s^{N-1} - (N-1)b_s^N$. The distribution function of Y_2 conditional on $Y_2 > b_s$, is

$$\Pr[Y_2 \leq y|Y_2 > b_s] = \frac{F_2(y) - F_2(b_s)}{1 - F_2(b_s)}, \text{ for any } y \in [b_s, 1]. \tag{4.22}$$

Differentiating (4.22) with respect to y, we obtain the probability density function f_{cond} of Y_2 conditional on $Y_2 > b_s$:

$$f_{cond}(y) = \frac{f_2(y)}{1 - F_2(y)} = \frac{N(N-1)(1-y)y^{N-2}}{1 - Nb_s^{N-1} + (N-1)b_s^N}. \tag{4.23}$$

Using this, the conditional expected value of Y_2 given that $Y_2 > b_s$ equals

$$E[Y_2|Y_2 > b_s] = \int_{b_s}^{1} y \cdot f_{cond}(y)\, dy = \frac{N(N-1)}{1 - Nb_s^{N-1} + (N-1)b_s^N} \cdot \left(\frac{1-b_s^N}{N} - \frac{1-b_s^{N+1}}{N+1} \right). \quad (4.24)$$

Combining (4.19), (4.20) and (4.21) with (4.24), the seller's expected profit is given by:

$$E_s(b_s) = b_s^N x_s + N(1-b_s)b_s^N + (N-1)(1-b_s^N) - N\frac{(N-1)}{(N+1)}(1-b_s^{N+1}). \quad (4.25)$$

In order to maximize E_s with respect to b_s, we differentiate $E_s(b_s)$.

$$\begin{aligned} E_s'(b_s) = 0 &\Leftrightarrow Nb_s^{N-1}x_s - Nb_s^N + N^2(1-b_s)b_s^{N-1} - (N-1)Nb_s^{N-1} + N(N-1)b_s^N = 0 \\ &\Leftrightarrow b_s = \beta_s^s(x_s) = \frac{x_s+1}{2}. \end{aligned} \quad (4.26)$$

■

Seller's expected profit in second-price auction. Replacing b_s with $\frac{x_s+1}{2}$ in Equation (4.25) we obtain the seller's optimal expected profit in a second-price auction in which the seller participates too.

$$E_{opt}^s(x_s) = \frac{(x_s+1)^{N+1}}{(N+1)2^N} + \frac{N-1}{N+1}. \quad (4.27)$$

We observe that the seller's optimal bid equals the optimal reserve price in the model of Section 4.1.1. Furthermore, the seller's optimal expected profit E_{opt}^s remains the same as in the case of imposing a reserve price. That is, it does *not* make any difference for the seller whether he announces the reserve price prior to the auction or not. Bidders' truthful bidding is not affected either.

As already discussed in 4.1.1 profit equivalence with first and second-price auctions does apply. It is apparent from (4.18) and (4.27), that profit equivalence does *not* apply with first and second-price auctions in case of seller participation. Next, we compare the respective expected profits.

4.3 Profits Comparison in First and Second-Price Auctions with Seller Participation

In Section 4.2 we restricted attention to bidding behavior given the type of the auction employed by the seller. In this setting, we derived the optimal expected profit for the seller in first and second-price auctions to be E_{opt}^f and E_{opt}^s respectively.

Since the optimal auction for selling a single good is the second-price auction with reserve price (see [25]), then the second-price auction with seller participation is an optimal auction too. Thus, one would expect that the following inequality holds for every $x_s \in [0,1]$: $E^s_{opt}(x_s) \geq E^f_{opt}(x_s)$. However, we prove below that the aforementioned inequality applies only for a subset of the range of values of x_s, as opposed to the entire interval [0,1]. In particular, we prove the following proposition:

Proposition 4.3.1 *Let* $\Delta E(x_s) = E^f_{opt}(x_s) - E^s_{opt}(x_s)$. *Then* $E^s_{opt}(x_s) > E^f_{opt}(x_s)$, *if* $x_s \in (a,b)$ *and* $E^s_{opt}(x_s) < E^f_{opt}(x_s)$, *if* ($x_s \in [0,a)$ *or* $x_s \in (b,1]$), *where* a *and* b *are the unique roots of* ΔE *in* $[0,1]$, *with* $a < b$.

Proof. Combining equations (4.18) and (4.27),

$$\begin{aligned}\Delta E(x_s) &= E^f_{opt}(x_s) - E^s_{opt}(x_s)\\ &= \frac{1}{N}\left(\frac{N+1}{N}\right)^{N-1} x_s^{N+1} + \left(\frac{N}{N+1}\right)^2 - \frac{(x_s+1)^{N+1}}{(N+1)2^N} - \frac{N-1}{N+1}.\end{aligned} \tag{4.28}$$

We take the first order derivative of ΔE with respect to x_s:

$$\begin{aligned}\Delta E'(x_s) = 0 &\Leftrightarrow \left(\frac{N+1}{N}\right)^N x_s^N - \frac{(x_s+1)^N}{2^N} = 0\\ &\Leftrightarrow x_s = \frac{N}{N+2}.\end{aligned} \tag{4.29}$$

The second order derivative of $\Delta E(x_s)$ equals:

$$\Delta E''(x_s) = N\left(\frac{N+1}{N}\right)^N x_s^{N-1} - \frac{N}{2}\frac{(x_s+1)^{N-1}}{2^{N-1}}. \tag{4.30}$$

It can be checked that

$$\begin{aligned}\Delta E''\left(\frac{N}{N+2}\right) &= \frac{(N+1)^N}{(N+2)^{N-1}} - \frac{N(N+1)}{2^{N-1}(N+2)}\\ &= \frac{N+1}{N+2}\left\{\frac{[2(N+1)]^{N-1} - N(N+2)^{N-2}}{2^{N-1}(N+2)^{N-2}}\right\} > 0.\end{aligned} \tag{4.31}$$

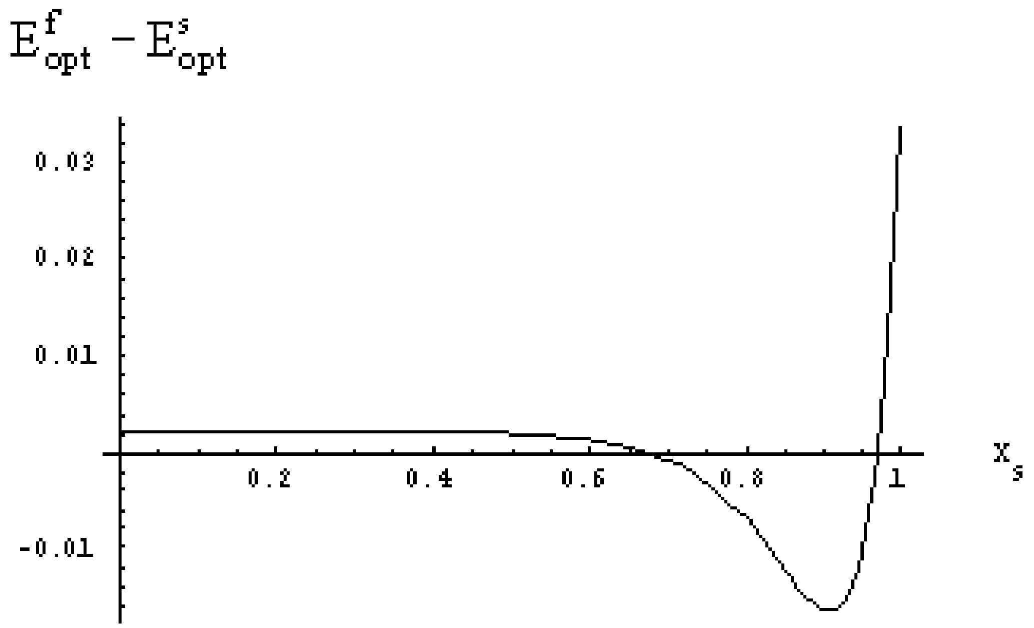

Figure 4.1: Profit comparison in first and second-price auctions with seller participation

Thus the function ΔE has a unique minimum point at $x_s = \frac{N}{N+2}$. Since

$$\Delta E(\frac{N}{N+2}) = \frac{(N+2)^N - (N+1)^{N+1}}{(N+2)^N(N+1)^2} < 0,$$

$$\Delta E(0) = \frac{(N+1)2^N - (N+1)^2}{(N+1)^2(N+1)2^N} > 0 \text{ and} \tag{4.32}$$

$$\Delta E(1) = \frac{(N+1)^{N+1} - (2N+1)N^N}{N^N(N+1)^2} > 0,$$

we conclude that ΔE has two roots in the interval [0,1]. Let a and b be those roots with $a < b$. For $N \geq 4$ these roots can only be found numerically. Then, $E^f_{opt} > E^s_{opt}$ if $x_s \in [0, a)$ or $x_s \in (b, 1]$ and $E^f_{opt} < E^s_{opt}$ if $x_s \in (a, b)$. ∎

Figure 4.1 illustrates the function $\Delta E(x_s)$ for $N = 20$. The roots of ΔE equal $a = 0.675$ and $b = 0.97$. The seller prefers a first-price auction if his valuation $x_s \in [0, 0.675) \cup (0.97, 1]$ and a second-price auction if $x_s \in (0.675, 0.97)$.

Note that the difference $\Delta E(x_s)$ converges to 0 as N tends to infinity. This means that the first and second-price auctions tend to yield the same expected profit to the seller in the presence of a very large set of bidders. Indeed, the bidders' shading decreases as the number of bidders increases. Actually, the bids converge to the bidders' true valuations. Thus, the expected payment under the first-price auction increases and tends to reach the maximum possible value which is 1. Similarly, the expected payment under the second-price auction increases with the number of bidders up to the value 1.

4.3.1 An Apparent Explanation Paradox and its Explanation

We have proved that the seller prefers a first-price auction in some cases and a second-price auction in others, depending on his own valuation. We also know from the literature, that a second-price auction with reserve price, or equivalently with seller participation, is an optimal auction with respect to seller's expected profit, independently of the seller's valuation. This may be seen as a paradox in the first place. Next, we argue that both results are valid because each applies in a different context.

The seller's choice regarding which of the two auctions to apply depends on his valuation. Each bidder knows that the seller chooses the auction that maximizes his expected profit. Thus, a bidder receives information about the seller's valuation. For example, consider that the seller adopts the strategy: choose first-price if $x_s \in [0, a) \cup (b, 1]$ or choose second-price if $x_s \in (a, b)$. If a particular seller runs a first (respectively second) price auction, then each of the bidders realizes that the seller's valuation x_s lies in $[0, a) \cup (b, 1]$ (respectively in (a, b)). The bidder's optimal strategy is now different than that derived in Section 4.2 for the particular mechanism. Indeed, there we did not take into consideration any extra information about the seller's valuation. We assumed that the seller is enforced to choose one type or the other and that his valuation is uniformly distributed on [0,1]. The results of Section 4.2 are valid only under this assumption.

In the general model in which the seller's choice of an auction type does affect bidding behavior, the set of equilibrium strategies includes the bidders' bidding strategies, the seller's bidding strategy and the distribution function of the seller's valuation. In order to verify that the expected profits of the first-price auction are indeed less than those of the optimal auction, the distribution function of the seller's valuation should be derived as part of the players' equilibrium strategies. The set of players' strategies of the second-price auction with seller participation we derived in Section 4.2, is an equilibrium of the general model, since bidders' strategy is truthful bidding irrespectively of the distribution function of the seller's valuation. This is indeed the optimal mechanism, thus our results do not contradict with those from literature.

Concluding Remarks

In this chapter we assess the impact of seller participation on bidding behavior and profits. We derive equilibrium strategies in first and second-price auctions and compare the expected profits obtained by the seller. This study can be extended to multi-unit auctions. In many markets such as bandwidth markets, the owner of the goods benefits by keeping a portion for himself. In this case he would require to obtain as much profit as possible as a seller, whereas he would benefit by an efficient selling procedure as a buyer. Analysis of such models constitutes an interesting yet challenging direction for future research.

Chapter 5

Bandwidth Allocation in a Communication Network: Assessment of PSP and a New Strategy

5.1 Introduction

We consider the problem of allocating a quantity C of bandwidth efficiently to a set of bidders in a communication *network* of arbitrary topology. Each bidder is assumed to have a privately known valuation. The feature that plays a significant role in designing the appropriate mechanism in this model, is the requirement that a bidder should obtain the *same* quantity of bandwidth in each link that belongs to his own path. (This is henceforth referred to as *consistent* allocation of bandwidth to this user). In the case where a bidder would obtain more quantity in a link than in the others, this excess quantity would be useless to him, while he would have paid for it a non-negative amount of money. This portion could be fairly given to another bidder increasing in this way both the social welfare and the net benefit of the two users. Therefore, consistent allocation of bandwidth is necessary for attaining efficiency. Consistent bandwidth allocation can be attained by means of a combinatorial auction mechanism that would allow a unique bid per bidder for the whole path. However, since in general users are interested in a multitude of different paths, the winner determination process of a combinatorial auction (described in Section 2.6) is very complicated, if not intractable. Lazar and Semret propose in [20] an auction mechanism that allows for independent bids in each link. Each bidder exploits the information of all links he acquires during the process and determines his own bid for each link independently. In the PSP auction a bidder is allowed to submit a bid in any link irrespectively of the network topology. The most common case is that of bidding in consecutive links that form a path but there can be other cases too, in which a bidder has interest in links that are not physically connected. In the latter case, the bidder following the PSP rules, submits a bid independently in each link of interest without being affected by the network topology. The aforementioned mechanism is a generalization of the Progressive Second Price (PSP) auction for a single link that we briefly described, in Chapter 2. The PSP auction has played a significant role in the field of auctioning network resources and has attracted much attention of other researchers too, who have published related works, e.g. [22, 23, 24, 32]. Prior to describing the network-wide PSP and the problems we have investigated, we give our remarks about the PSP auction for a single link.

5.2 Remarks on the Progressive Second Price Auction for a Single Link

The PSP auction manages to allocate a divisible good such as bandwidth in a nearly efficient way. Nevertheless, there have been observed some drawbacks of the mechanism which we briefly discuss below. A major drawback of PSP, is that the strategy of the ϵ-best reply is considered in the short run. That is, without taking into account the future actions of opponent bidders in the determination of one's optimal strategy. Thus, the ϵ-best reply is ϵ-optimal (i.e., nearly optimal) for a bidder, if this bid would cause the termination of the game, as opposed to the repeated game. As evidence to this, Maillé and Tuffin provide in [22] a different Nash equilibrium of the PSP game, which comprises strategies that take into consideration subsequent actions: the first bidder observes no demand by his opponents. Thus, he has the incentive to bid for the whole capacity at a very high price so that no other bidder can afford bidding against him afterwards. As a result the first bidder wins all the capacity at a very low price (reserve price p_0). Extending this argument, one could claim that a bidder at any possible step of the PSP game could *overbid* for the desired quantity q^ϵ to exclude all subsequent bidders. However, in our opinion, this does not apply. Indeed, a bidder cannot know or determine whether he will be playing first or at any specific order, since the game is completely asynchronous and distributed. This implies that more than one bidders might submit a bid, each having the belief that he is the first to play. In this case, overbidding might cause a bidder to suffer a negative net benefit. This follows by reasoning as in the case of the Vickrey auction, that a bidder cannot profit from overbidding: he faces the risk to obtain the good at a price higher than his own valuation. The difference is that overbidding is never preferred in the Vickrey auction whereas in the PSP auction the bidder might be better off by overbidding, but he cannot predict when. The following example illustrates the above ideas.

Example 5.2.1

Suppose that two bidders A and B compete for a quantity $C = 5$. Bidders A and B have marginal valuations $\theta'_A(x) = -x+6$ and $\theta'_B(x) = -x+9$ respectively. Assume that each bidder plays as if he were the first one. Let $s_A = (5, 200)$ and $s_B = (5, 100)$ be their bids. Each one overbids expecting to exclude the other from participating. The winner is bidder A who obtains the whole capacity at a unit price of 100. Thus, his net benefit equals $NB_A = \theta_A(5) - c_A(5) = -5^2/2 + 6 * 5 - 500 = -482.5$. ▲

In our opinion, the ϵ-Nash equilibrium established by Lazar and Semret in the single-link case is a reasonable set of strategies for the bidders. Nevertheless, Maillé and Tuffin propose in [22] a variation of the PSP algorithm that deters the first bidder from overbidding. According to the authors of [22] bidders that cannot get any bandwidth are required to submit a bid, at no cost, which will punish the bidder who obtains all the capacity due to overbidding.

Our second remark has to do with the case of *ties*. That is, when a bidder's bidding price coincides with another previously submitted bidding price. Then, the PSP allocation rule punishes both bidders involved. For example, assume two bidders that compete for capacity $C = 8$ and submit the bids $s_1 = (4, 7)$ and $s_2 = (6, 7)$ respectively. According to the PSP allocation rule bidder 1 obtains $8 - 6 = 2$ units of bandwidth and bidder 2 obtains $8 - 4 = 4$ units of bandwidth. The remaining 2 units of bandwidth are not allocated to them. This perspective gives bidders the incentive to change their bids by increasing slightly their price and decreasing the demanded quantity, so as to avoid ties. As deduced from [20], the intuition for introducing ϵ is that bidders never experience ties if they follow the strategy of the ϵ-best reply explained in the previous chapter. However, this is *not* true; any selected price near the best reply might coincide with another bid, although the introduction of ϵ does reduce the frequency with which ties occur. In our opinion, this is actually the impact of parameter ϵ, as opposed to the interpretation given in [20], that ϵ can be thought of as fee each time a bidder submits a bid. With this approach, a bidder faces the risk of obtaining a negative net benefit, since the strategy he adopts is based only in the current step. In the extreme case, a bidder that places K bids and finally wins no bandwidth at all, faces a negative net benefit of $-K\epsilon$ units. Tuffin considers in [32], the problem of ties. He defines a new allocation rule and proves that ties do not occur, in the PSP auction under this new rule. This modified allocation rule reduces to the rule in equation (2.3) when no tie arises. Thus the strategy proposed by Lazar and Semret in [19] and the properties of convergence to an ϵ-Nash equilibrium and efficiency still apply.

5.3 PSP in a Network

Lazar and Semret extend in [20], the PSP auction in the network case. They consider a set of bidders $\mathcal{N} = \{1, \ldots, N\}$ and a set of resources $\mathcal{L} = \{1, \ldots, L\}$ corresponding to communication links of a network, with bandwidth capacity $C^1, \ldots, C^L$ respectively. Each bidder desires bandwidth in a combination of links that form his path, which is fixed and known to him before the auction starts. Bids at any combination of links are allowed. All the assumptions about bidders' valuations made in the single link case still hold. It is further assumed that each bidder benefits only from the *minimum* allocated quantity of bandwidth among all links of his path. For example, consider a bidder interested in a path consisting of two links. Consider further, that he obtains 3 and 5 units of bandwidth in each link respectively. He can use only 3 units in both links, even though he has paid for more in the second link. The PSP auction is performed in each link independently, so that the whole mechanism is decentralized: the outcome at any link can be extracted without knowledge of other links' state. Each bidder combines information released from all links he is interested in, and determines his strategy such that his net benefit be maximized.

Lazar and Semret in [20], prove that a bidder cannot do better than place the *same* bid at all links on his path (consistent bidding). They claim that there exists a truthful ϵ-best reply, which leads the game at a truthful ϵ-Nash equilibrium if reserve prices for each link are introduced. The

ϵ-best reply $s^\epsilon = (q_i^\epsilon, p_i^\epsilon)$ of bidder i that is submitted in each link separately is derived as follows: the best reply $s = (q_i, p_i)$ is the intersection of i's marginal valuation function and the "staircase" that depicts the market price at each quantity. For bidder i, the market price at quantity q equals the *sum* of the prices at quantity q offered by the opponents at the auctions of the various links that form bidder's i path. That is, the "market staircase" of bidder i equals the sum of the "staircases" of each link in his path: for each quantity are added prices. In other words, q_i is the largest quantity such that his marginal value is just greater than the market price, whereas p_i equals his marginal valuation at q_i. That is, $p_i = \theta_i'(q_i)$. The ϵ-best reply is then calculated by decreasing q_i by $\frac{\epsilon}{\theta'(0)}$ so that $q_i^\epsilon = q_i - \frac{\epsilon}{\theta'(0)}$, and adjusting p_i such that $p_i^\epsilon = \theta_i'(q_i^\epsilon)$. The bidder then bids $(q_i^\epsilon, p_i^\epsilon)$ in each link of his path. Figure 5.1 illustrates the derivation of the ϵ-best reply of bidder i whose path comprises two links with the same capacity, given his opponents' profiles in the two links.

Figure 5.2 illustrates the derivation of the ϵ-best reply of bidder i whose path comprises two links with different capacities C_1 and C_2 ($C_1 < C_2$) respectively, given his opponents' profiles in the two links. In the case where $C_1 \neq C_2$, bidder's i demand function θ_i' is defined in the interval $[0, \min\{C_1, C_2\}]$, since he requires to obtain the same quantity in each link. Thus, the monotonicity property of his bid (price increases and quantity decreases in time) is preserved: the bid never lies in the interval $[\min\{C_1, C_2\}, \max\{C_1, C_2\}]$. For simplicity reasons, we consider links with the same capacity in the rest of this chapter.

Efficiency in the network-wide PSP auction is examined in [31]. It is proved in [31] that the social welfare is within a bound from its maximum value provided that the second derivatives of bidders' valuation functions are bounded.

Efficiency and related aspects in the network-wide PSP auction. The aforementioned bidding strategy for the network case suggests that path bidders submit a high bid, the same in each link, for the demanded quantity. Moreover, in this step, they could obtain this same quantity by submitting a lower price in each link. This overbidding is apparently harmless to the path bidder himself, due to the fact that charging is performed according to social opportunity cost. By construction of the market "staircase", it follows that the total charge in all links never exceeds p_i^ϵ. Nevertheless, we claim that efficiency is not very close to maximum in many cases, due to this overbidding approach. Competitive path bidders are reasonably expected to have higher

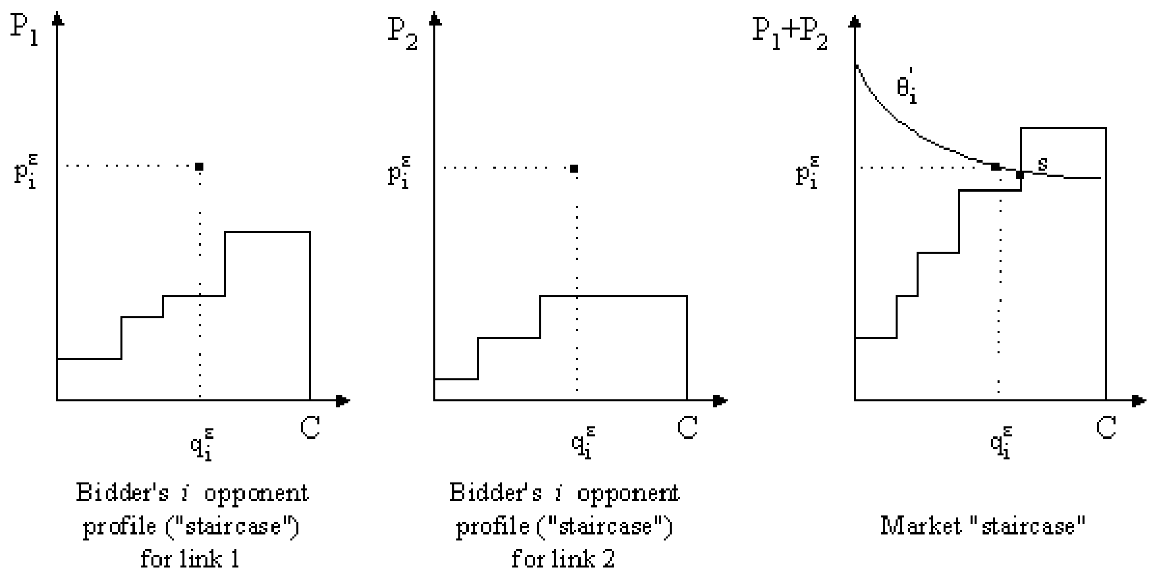

Figure 5.1: Bidder's i ϵ-best reply $(q_i^\epsilon, p_i^\epsilon)$ proposed in [20] for a network of two links with the same capacity

valuations in general than single-link bidders, because their valuations correspond to a whole path instead of one single link, even if single-link bidders should win considerable portions of capacity in the efficient allocation. In general, the path bidder who desires the path with the most links, has the advantage over all other bidders to submit a high bid and exclude them from winning. This problem is reminiscent of the threshold problem discussed in Chapter 2: combinatorial bidding causes a low valuation bidder to be unable to respond to high bids. Bidding in network-wide PSP is similar to that of combinatorial, since each bidder's strategy is derived by considering the sum of the market prices of all links as opposed to the set of the different prices of each link. We have carried out several experiments that verify our assertion. These will be described in a subsequent section. Herein, we provide an example in which the approach discussed leads to an inefficient allocation.

Example 5.3.1

Consider two communication links A and B with capacities $C^A = C^B = 8$ and three bidders with marginal valuations $\theta_1'(x) = -x + 10$ for link A, $\theta_2'(x) = -x + 16$ for link B, $\theta_3'(x) = -x + 25$ for

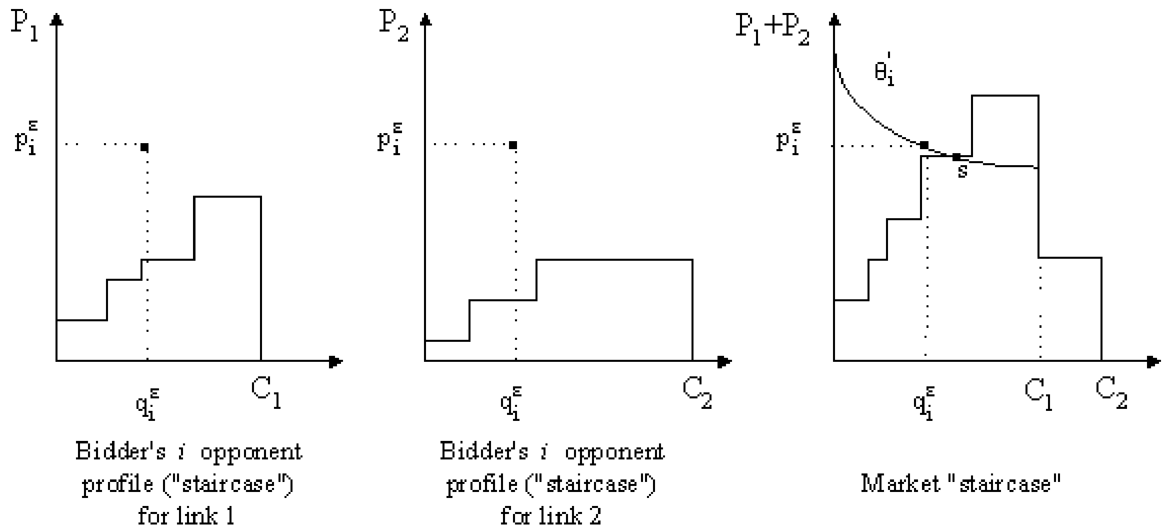

Figure 5.2: Bidder's i ϵ-best reply $(q_i^\epsilon, p_i^\epsilon)$ proposed in [20] for a network of two links with different capacities

links A and B respectively. Their respective utilities are $\theta_1(x) = -\frac{x^2}{2} + 10x$, $\theta_2(x) = -\frac{x^2}{2} + 16x$ and $\theta_3(x) = -\frac{x^2}{2} + 25x$. We take ϵ equal to 1 and the reserve price p_0 in each link equal to 1 too. The efficient allocation in this example is the solution of the following social welfare maximization problem:

$$\begin{aligned} \max_{\vec{x}} \; \{\theta_1(x_1) + \theta_2(x_2) + \theta_3(x_3)\} \Leftrightarrow \max_{\vec{x}}\{-\frac{x_1^2}{2} + 10x_1 - \frac{x_2^2}{2} + 16x_2 - \frac{x_3^2}{2} + 25x_3\} \\ \text{s.t. } x_1 + x_3 = 8 \\ x_2 + x_3 = 8 \end{aligned} \quad (5.1)$$

The solution of this problem is the vector of quantities $\vec{x} = (3, 3, 5)$ and the maximum value of the social welfare is $SW_{opt}(3, 3, 5) = 181.5$.

If we apply the network-wide PSP auction and the proposed strategies for the bidders, then the game terminates immediately after the first bid of the path bidder 3, independently of the order the bidders may submit their bids. For example, assume for the sake of simplicity that bidder 3 plays first; that is, he observes no demand in neither of the two links. His bid $s_3^\epsilon = (q_3^\epsilon, p_3^\epsilon)$ for each link according to the strategy is calculated as follows: $q_3^\epsilon = 8 - \frac{\epsilon}{\theta_3'(0)} = 7.96$ and $p_3^\epsilon = \theta_3'(7.96) = 17.04$.

The other single-link bidders cannot compete against bidder 3, since the marginal valuation of any one of them is less than 17.04 for any quantity of bandwidth. We assume that the remaining quantity of bandwidth which equals 0.04 units in each link is awarded to the respective single-link bidder at the reserve price. (This would be the case if the single-link bidders played first.) Thus, the auction terminates at the allocation $\vec{x} = (0.4, 0.4, 7.96)$. The resulting social welfare $SW(0.4, 0.4, 7.96)$ equals 168.3576, which is significantly below the optimal value. ▲

Next, we justify why the social welfare of the network-wide PSP auction often deviates considerably from the corresponding optimal value. This is due to the fact that the ϵ-best reply is *not truthful* in the following sense: The ϵ-best reply is a point of the overall marginal valuation function of bidder i. However, on a per link basis, the ϵ-best reply submitted is not a truthful one, since the sum of the prices p_i^ϵ offered for the quantity q_i^ϵ in each link is higher than the bidder's marginal valuation at q_i^ϵ which is p_i^ϵ. Actually, the price p_i^ϵ offered in each link is considerably higher than the price that would suffice to obtain the quantity q_i^ϵ. Since the auction in each link is a PSP auction itself, this strategy amounts to overbidding. However, overbidding forces prices increase more aggressive in each link, which thus does not converge to those resulting in the efficient outcome, as was the case in the single-link PSP auction. Next, we provide an example in which the deviation of the social welfare from its maximum value can be very high as the number of links increases.

Example 5.3.2

Consider a linear network consisting of L links, each of capacity C. Consider further that there is one path bidder interested for the whole path of the L links and one single-link bidder in each link. We take that the utility function θ_p of the path bidder and the utility functions θ_i of the various single-link bidders are identical. Let θ be the common utility function of all bidders. Then, the efficient allocation is to award each single-link bidder with capacity C resulting in the optimal social welfare $SW_{opt} = L\theta(C)$. If each bidder adopts the strategy of Lazar and Semret and the path bidder plays first, then the path bidder obtains the whole capacity due to overbidding. The resulting social welfare equals $SW = \theta(C)$. Thus, the percentage loss of the social welfare equals $loss = \frac{L\theta(C) - \theta(C)}{L\theta(C)} = \frac{L-1}{L}$. Note that the loss increases and approaches 100% as the number of links increases. ▲

5.4 A New Bidding Strategy for Path Bidders

As already mentioned, a key question on the proposed strategy for path bidders of the network-wide PSP auction is whether it is fair or optimal (for the bidder himself and/or for the society) to apply overbidding in the various links. It is not obvious whether this approach leads to optimal price discovery and final bandwidth allocation. Thus, we propose another strategy for path bidders of the network-wide PSP auction that we consider to be more reasonable. This strategy suggests that instead of bidding the market clearing price in each link, split this and set the prices -differently in each link- to be the smallest ones that assure winning of the optimal quantity. The optimal quantity is found similarly as in the strategy of [20].

Recall that, by assumption, each bidder knows his valuation function for the whole path of his interest. If he could split his valuation function and determine the corresponding portion for each link, then he could find the right price for this link's bid. For each link, this price corresponds to the intersection of the vertical line through the optimal quantity with *this* link's "staircase". These prices sum to the market clearing price. Therefore, a path bidder employing this strategy reveals his real demand, thus *avoiding* overbidding.

Formally, consider a bidder i that is interested in a path consisting of K links, which are taken to be links $1, \ldots, K$. Let $s = (q_i, p_i)$ and $s^\epsilon = (q_i^\epsilon, p_i^\epsilon)$ be the best reply and the ϵ-best reply respectively, as defined in the previous section. For $j = 1, \ldots, K$, let I_j be the set of all intersection points of the vertical line through q_i with the "staircase" of link j. For the set of points I_j there are two possibilities:

a) I_j is a single point that lies on the interior of a horizontal segment of the "staircase" of link j; this is the general case.

b) I_j is a set of points that lie on a vertical segment of the "staircase" of link j. This means that the quantity q_i coincides with another previously submitted quantity in link j. This is an exceptional case arising when two or more bidders are identical.

In order to define the price $p_{i,j}$ that bidder i will submit in each link j, we distinguish among the following three cases:

Case 1. For each $j = 1, \ldots, K$, I_j is a single point. Then, define $p_{i,j}$ to be the price that corresponds

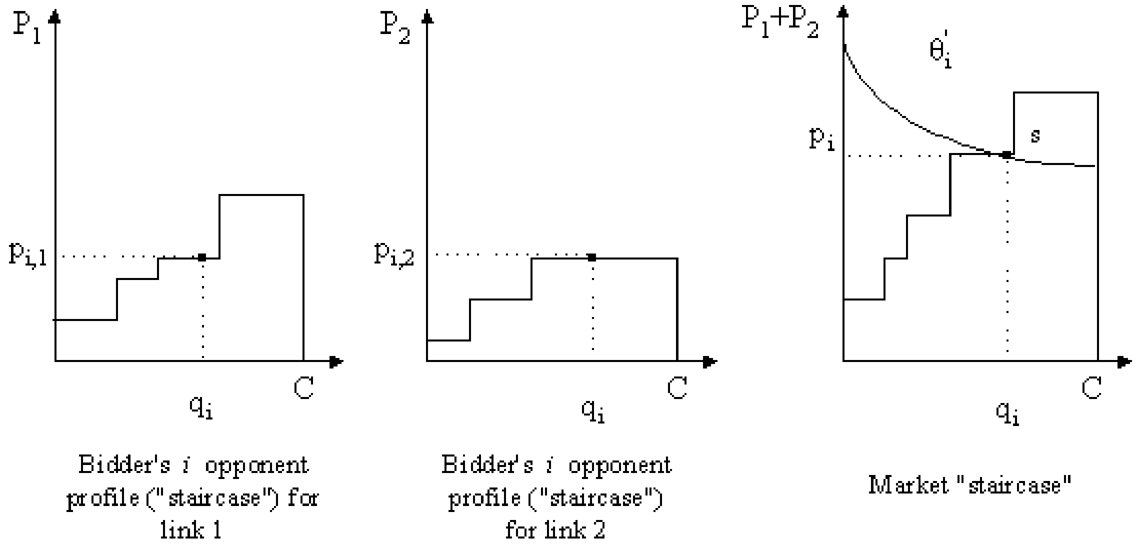

Figure 5.3: The vertical line through q_i intersects with the "staircase" of both links horizontally (Case 1)

to the intersection point I_j. Obviously, prices $p_{i,j}$ are uniquely defined in this case. Figure 5.3, depicts the prices $p_{i,1}$ and $p_{i,2}$ for bidder i in a two-link path in this case.

Case 2. For each $j = 1, \ldots, K$, $j \neq k$, I_j is a single point. I_k is a set of points. Then, for each $j = 1, \ldots, K$, $j \neq k$ define $p_{i,j}$ as previously and let $p_{i,k}$ equal $\theta'(q_i) - (\sum_{j \neq k} p_{i,j})$. Again, prices $p_{i,j}$ are uniquely defined. Figure 5.4 depicts the prices $p_{i,1}$ and $p_{i,2}$ in a two-link path in this case.

Case 3. For two or more links, the respective I_j's are intervals rather than single points. Then, for each such link j, define $p_{i,j}$ to be the lowest price of all prices that correspond to the intersection points of set I_j. For the remaining links define $p_{i,j}$ as in case 1. Prices $p_{i,j}$ are uniquely defined in this case too. Figure 5.5 depicts prices $p_{i,1}$ and $p_{i,2}$ in a two-link path, in this case.

Recall that the market price p_i at quantity q_i equals the sum of the prices at quantity q_i offered by i's opponents at the auctions of the various links that form bidder's i path. From the definition

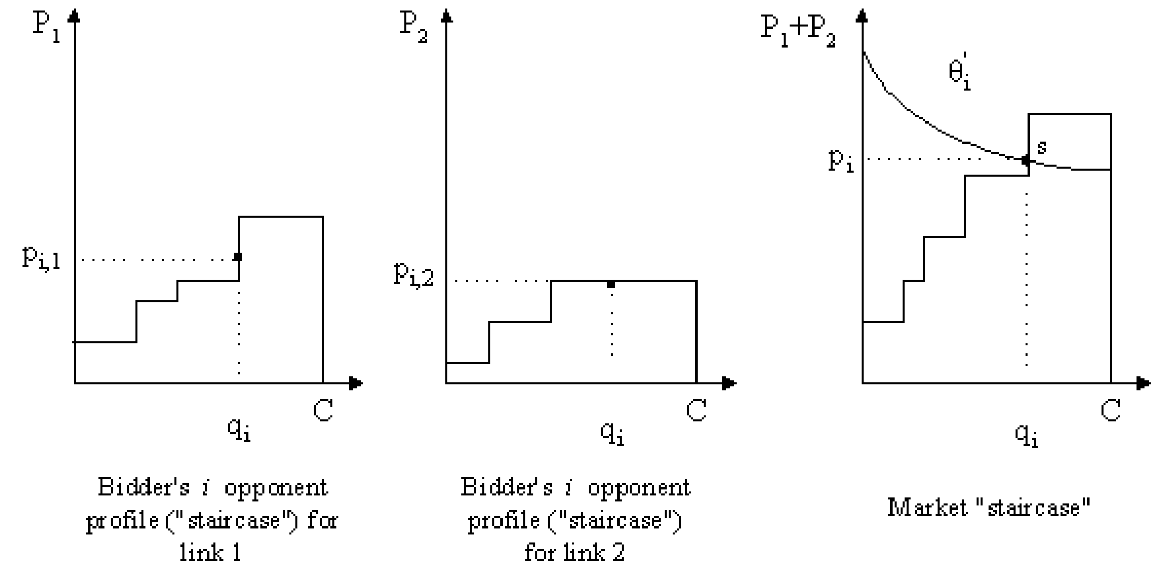

Figure 5.4: The vertical line through q_i intersects with the "staircase" of link 1 vertically and with the "staircase" of link 2 horizontally (Case 2)

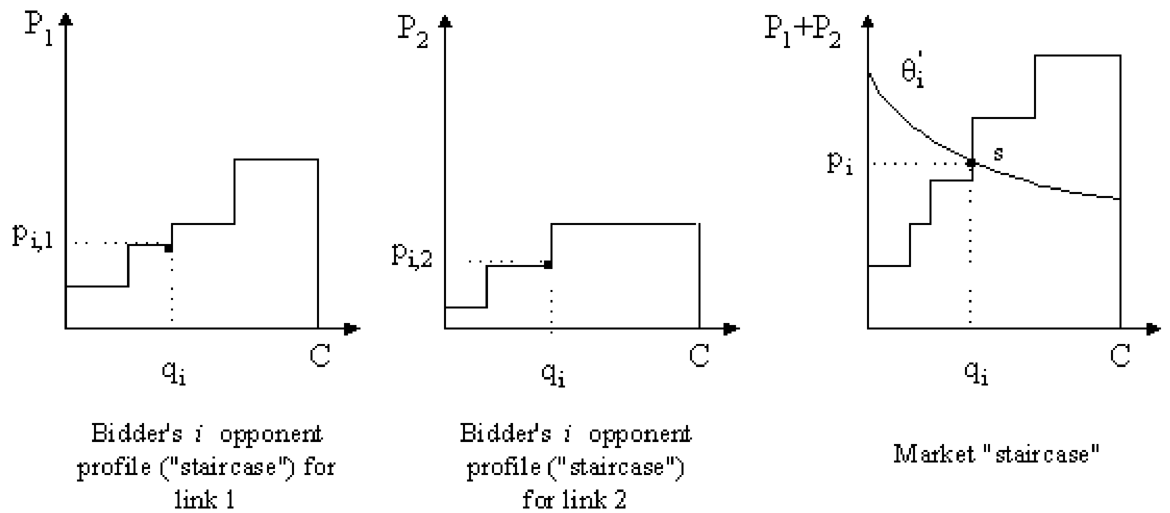

Figure 5.5: The vertical line through q_i intersects with the "staircase" of both links vertically (Case3)

of $p_{i,j}$'s we have the following:

$$p_{i,1} + \cdots + p_{i,K} = p_i, \text{ for cases 1 and 2 and,} \tag{5.2}$$

$$p_{i,1} + \cdots + p_{i,K} \leq p_i, \text{ for case 3} \tag{5.3}$$

Similarly with the strategy of [20], the quantity to be demanded in the bids of all links is slightly less than q_i. To keep in line with [20] as much as possible, this decrease equals $\frac{\epsilon}{\theta_i'(0)}$. Since price increases as quantity decreases, we have that $p_i^\epsilon = p_i + G$, where p_i^ϵ is such that $p_i^\epsilon = \theta_i'(q_i^\epsilon)$ and G is some positive scalar. The prices of the bids of the various links should also be increased. For simplicity, we define

$$p_{i,j}^\epsilon = p_{i,j} + \frac{G}{K}.$$

(Alternatively, G could have been split proportionally with respect to the $p_{i,j}$.) Combining this with equations 5.2 and 5.3 we obtain the following:

$$p_{i,1}^\epsilon + \cdots + p_{i,K}^\epsilon = p_{i,1} + \cdots + p_{i,K} + G = p_i + G = p_i^\epsilon, \text{ for cases 1 and 2 and,} \tag{5.4}$$

$$p_{i,1}^\epsilon + \cdots + p_{i,K}^\epsilon \leq p_{i,1} + \cdots + p_{i,K} + G \leq p_i + G = p_i^\epsilon, \text{ for case 3.} \tag{5.5}$$

Bidder's i bid for link j is $(q_i^\epsilon, p_{i,j}^\epsilon)$ for $j = 1, \ldots, K$. Essentially, path bidder i first derives the best quantity and (total) price for the entire path and then *splits the price* among the various links so that he can win the desired quantity in all links. Obviously, for single-link bidders, this strategy reduces to the one proposed by Lazar and Semret in [20] in the single-link case. Figure 5.6 depicts the derivation of bidder's i bids according to our approach for a two-link network.

Lazar and Semret have showed in [20], that the strategy they propose is the ϵ-best reply for each bidder in the *short run*; that is, without taking into consideration subsequent actions of his opponents. Our strategy is an ϵ-best reply in the short run too in the following sense: given a particular “staircase”, if the path bidder responds according to our strategy, then he obtains the same quantity and pays the same total charge as he would with the overbid strategy. Indeed, the quantity bidded for in our strategy is the same and even though the price in each link is lower than that of the overbid strategy, it is high enough so that the desired quantity be won. Since the charge in each link is the social opportunity cost, it is the same in both cases had the auction terminated with this bid. Thus, the path bidder's net benefit is the same with that under the overbid strategy, which would be optimal, had this bid caused termination of the auction.

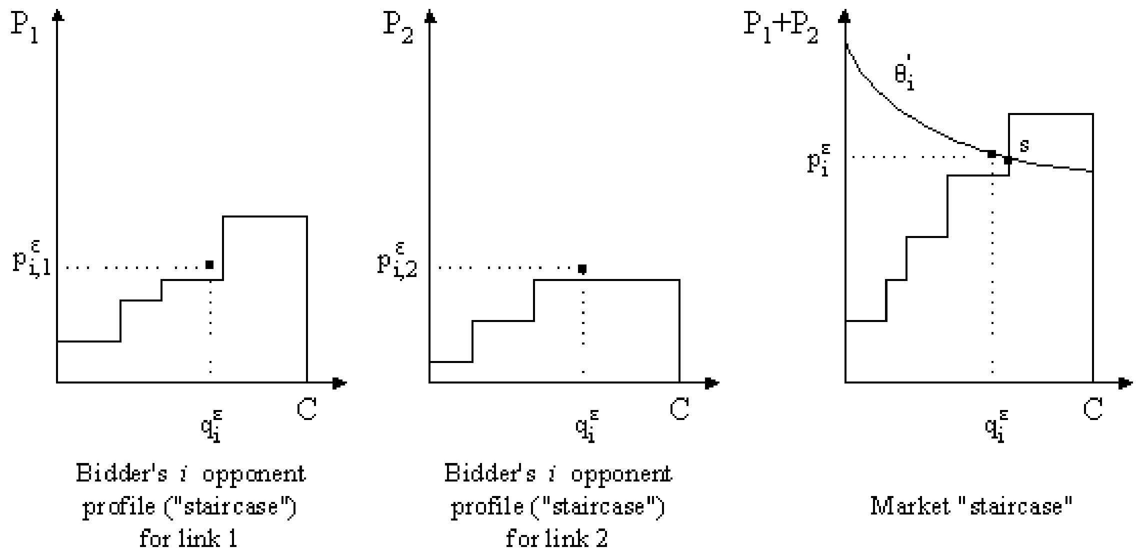

Figure 5.6: Bidder i submits the bid (q_i^ϵ, p_i^1) in link 1 and (q_i^ϵ, p_i^2) in link 2 according to the proposed strategy in a network of two links

However, the two strategies differ in how the auction evolves in subsequent steps and of course in the final allocation.

Below, we compare the strategy of Lazar and Semret (overbid strategy, also denoted as "LS") with our proposed strategy (minimum bid strategy, also denoted as "new") in terms of both individual net benefit for path bidders and social welfare.

5.5 Experimental Evaluation of Strategies

In this section, we compare the overbid strategy with the minimum bid strategy in terms of efficiency and equilibrium. We expect that, since the minimum bid strategy amounts to less aggressive bidding, it yields outcomes that are closer to the maximum social welfare. It is also reasonable that path bidders with high valuation prefer the minimum bid strategy, since they will win anyway their share (as in efficient allocation), but without increasing prices that would result in a higher payment and a lower net benefit for them. On the contrary, path bidders (particularly those with low valuation) may prefer the overbid strategy, since they win more capacity than they obtain under the efficient allocation. These assertions are justified in the experiments that follow.

5.5.1 Specification of Experiments

We carried out several experiments and considered the following assumptions:

- The network consists of 2 links, each having capacity $C = 5$ units of bandwidth. Motivation for considering such a small network is provided at the end of this section.

- Each experiment involves two runs of the network-wide PSP auction. In the first run the overbid strategy is applied by every bidder. In the second run the minimum bid strategy is applied by every bidder. Additionally, we have carried out several experiments in which one bidder applies one of the two strategies and all the other bidders apply the other strategy.

- Players have elastic demand of the form $\theta'(q) = aq + b,\ a < 0$. Three categories of bidders take part in the game: single-link bidders interested in link 1, single-link bidders interested in link 2 and path bidders. We carried out experiments with the following mixes of bidders: a) many single-link bidders and one path bidder, b) many single-link bidders and a few path bidders, and c) a few single-link bidders and many path bidders.

- We consider a reserve price $p_0 = 0.1$ in each link. This price is set by the seller (bidder 0) of the first link and by the seller (bidder 1) of the second link.

- We take ϵ equal 1. When applying the minimum bid strategy for path bidder i, we split the extra price $\frac{\epsilon}{\theta_i'(0)}$ equally for the two links, so that $p_{i,1}^{\epsilon} = p_{i,1} + 0.5\frac{\epsilon}{\theta_i'(0)}$ and $p_{i,2}^{\epsilon} = p_{i,2} + 0.5\frac{\epsilon}{\theta_i'(0)}$, (We employ the notation defined in Section 5.4).

- The order of bidders is predefined and is applied periodically throughout an experiment.

Note that the PSP auction algorithm that we implemented differs from that of [20] in two points. First, since the problem of ties does not completely disappear by introducing ϵ (see Section 5.2), we introduced an iterative method to *avoid ties*. This is as follows: the ϵ-best reply is calculated, as in [20]. If a tie arises, then the price is reduced again by the same value as the first reduction; that is, by $\epsilon/\theta_i'(0)$. The quantity is further adjusted so that $p_i' = \theta_i'(q_i')$ is satisfied for the new bid (q_i', p_i'). We repeat this adjustment procedure until the tie is resolved. Second, we do *not* interpret ϵ as a bid fee. That is, a bidder submits his next bid if he thus improves his net benefit by any positive amount. Recall that in [20] it is necessary that this improvement is at least ϵ; otherwise, the bid is not submitted. Of course, ϵ is used to derive the ϵ-best reply, which is further employed to calculate the next value of net benefit. In our opinion, there is no reason to include ϵ again in the criterion of adjusting one's net benefit. Therefore, the algorithm bidder i employs is as follows:

1. Initially $s_i = 0$.
2. Take the updated s_{-i}.
3. Compute the ϵ-best reply s_i^ϵ as given by the corresponding strategy.
4. If necessary, adjust the bid s_i^ϵ so as to eliminate ties.
5. Let s_f be the bid at the end of step 3 and 4 (if necessary). If $NB_i(s_f) > NB_i(s_i)$, then replace the bid s_i with s_f.
6. In the next time of bidding go to step 2.

For the experimental evaluation of the two strategies we have developed a special software using the Java programming language. The optimal value of social welfare in each experiment was derived using Mathematica, considering full information about bidders' utility functions.

5.5.2 Results

The experiments confirmed our assertions for efficiency and bidding behavior in the network-wide PSP auction. In particular, the experiments revealed that the minimum bid strategy outperforms the overbid strategy with respect to efficiency of the mechanism, while it is also beneficial for path bidders. We will discuss about these findings in further detail below. Table 5.1 shows a typical representation of a pair of experiments. Each row provides information and results for a specific bidder. The first column shows the identity of each bidder, the second column refers to the link(s) each bidder is interested in, and columns 3 and 4 indicate the parameters a and b of the marginal valuation of each bidder. The optimal allocation is indicated in column 5. Each bidder's allocation and net benefit under the overbid strategy is indicated in column 6. Each

Bidder	**Link**	a	b	**Optimal Alloc**	**Over. Strategy**		**New strategy**	
					Bandw.	NB	Bandw.	NB
0 (seller)	1	0	0.1	0	0	0	0	0
2	1	-1	13	2.2	0	0	2.13	24.54
4	1	-1	12	1.2	0	0	1.19	12.79
6	1	-1	11	0.2	0	0	0.18	1.3182
8	1	-1	9	0	0	0	0	0
10	1	-1	10	0	0.05	0.09	0.05	0.1
1 (seller)	2	0	0.1	0	0	0	0	0
3	2	-1	3	0	0	0	0	0
5	2	-1	6	0.2	0	0	0.15	0.29
7	2	-1	7	1.2	0	0	1.18	6.49
9	2	-1	8	2.2	0	0	2.2	13.65
11	2	-1	5.5	0	0.05	0.13	0.02	0.13
12	1-2	-1	18	1.4	4.94	19.96	1.42	22.63
SW				90.3	77.63		90.24	

Table 5.1: A pair of experiments that shows the allocation and net benefit of each bidder in both strategies

bidder's allocation and net benefit under the minimum bid strategy is indicated in column 7. The last row shows the optimal social welfare, the social welfare under the overbid strategy and the social welfare under the minimum bid strategy.

Assessment of efficiency in the network-wide PSP auction. Running the same experiment with both strategies, the overbid strategy results in *lower* social welfare than with the minimum bid strategy. This was the case for *every* experiment we carried out. In fact, the more path bidders involved or the higher their valuation, the greater the observed loss in social welfare with the overbid strategy. Below we present the results of five sets of experiments with different mixes of path and single-link bidders. For each set of experiments, a variety of bidders' valuations were taken.

- Set A: 6 single-link bidders contend in link 1, 6 single-link bidders contend in link 2, while there is also 1 path bidder (interested in both links). In Table 5.2, we present the optimal social welfare, the social welfare under the overbid strategy and the social welfare under the minimum bid strategy for a multitude of experiments. The two last columns show the percentage loss of social welfare under the overbid and the minimum bid strategy respectively.
- Set B: 7 single-link bidders contend in link 1, 5 single-link bidders contend in link 2, while there is also 1 path bidder (interested in both links). In Table 5.3, we present the optimal social welfare, the social welfare under the overbid strategy and the social welfare under the minimum bid strategy for a multitude of experiments. The two last columns show the percentage loss of social welfare under the overbid and the minimum bid strategy respectively.
- Set C: 7 single-link bidders contend in link 1, 5 single-link bidders contend in link 2, while there are also 2 path bidders (interested in both links). In Table 5.4, we present the optimal social welfare, the social welfare under the overbid strategy and the social welfare under the minimum bid strategy for a multitude of experiments. The two last columns show the percentage loss of social welfare under the overbid and the minimum bid strategy respectively.
- Set D: 4 single-link bidders contend in link 1, 4 single-link bidders contend in link 2, while there are also 4 path bidders (interested in both links). In Table 5.5, we present the optimal social welfare, the social welfare under the overbid strategy and the social welfare under the minimum bid strategy for a multitude of experiments. The two last columns show the percentage loss of social welfare under the overbid and the minimum bid strategy respectively.
- Set E: 2 single-link bidders contend in link 1, 2 single-link bidders contend in link 2, while there are also 6 path bidders (interested in both links). In Table 5.6, we present the optimal social welfare, the social welfare under the overbid strategy and the social welfare under the minimum bid strategy for a multitude of experiments. The two last columns show the percentage loss of social welfare under the overbid and the minimum bid strategy respectively.

Experiment	Optimal SW	SW_{over}	SW_{new}	$loss_{over}(\%)$	$loss_{new}(\%)$
1	187.5	186.97	187.12	0.28	0.2
2	137.5	137.28	137.39	0.16	0.08
3	110.58	109.47	110.58	1.003	0
4	90.3	77.63	90.24	14.03	0.06
5	98.45	96.92	98.44	1.55	0.01
6	247.5	247.36	247.49	0.05	0.004
7	337.5	337.14	337.24	0.1	0.07
8	493.81	487.53	493.36	1.27	0.09

Table 5.2: Social welfare comparison in set A of experiments

The results of Tables 5.2 - 5.6, reveal the following: In each experiment the minimum bid strategy yields higher or equal social welfare than the overbid strategy. Moreover, social welfare under the minimum bid strategy is always almost equal to the optimal. The maximum social welfare loss observed was less than 0.22%. On the contrary, the social welfare under the overbid strategy may considerably deviate from the optimal value. For example, in the experiment 2 of set E the loss observed was 31.82%, in the experiment 5 of set D the loss observed was 8.94% and in the experiment 2 of set B the loss observed was 9.72%. In fact, the higher the number of path bidders considered, the higher the social welfare loss in the overbid strategy is observed. In the experiments of set E, the minimum loss of social welfare is 4.44% and the maximum loss of social welfare is 31.82% under the overbid strategy. In most experiments of set A, the social welfare loss under the overbid strategy is very close to that under the minimum bid strategy, since there is only one path bidder. Nevertheless, in experiment 4 of set A, the loss under the overbid strategy is remarkably high 14.03% because the path bidder in this experiment has a high valuation and this overbidding is beneficial for him. Experiments 5,6,7 of set E comprise the same set of users, but in a different order. We observe that the social welfare under the minimum bid strategy is the same for all these experiments, while the social welfare under the overbid strategy differs in each of these experiments.

Experiment	Optimal SW	SW_{over}	SW_{new}	$loss_{over}(\%)$	$loss_{new}(\%)$
1	215.67	193.73	215.60	10.17	0.03
2	215.70	194.73	215.64	9.72	0.02
3	215.76	195.72	215.68	9.28	0.03
4	215.85	196.71	215.77	8.86	0.03
5	215.96	197.71	215.87	8.45	0.04
6	220.30	212.63	219.99	3.48	0.14
7	237.70	237.51	237.70	0.07	0
8	262.5	262.42	262.42	0.03	0.03

Table 5.3: Social welfare comparison in set B of experiments

Experiment	Optimal SW	SW_{over}	SW_{new}	$loss_{over}(\%)$	$loss_{new}(\%)$
1	262.5	262.39	262.42	0.04	0.03
2	237.75	216.62	237.48	8.88	0.11
3	223.09	207.65	222.97	6.92	0.05
4	223.95	212.63	223.75	5.05	0.08
5	225.45	217.60	225.15	3.48	0.13
6	243.75	237.51	243.74	2.56	0.004
7	233.36	232.54	233.33	0.35	0.01
8	247.5	247.48	247.48	0.008	0.008

Table 5.4: Social welfare comparison in set C of experiments

Experiment	Optimal SW	SW_{over}	SW_{new}	$loss_{over}(\%)$	$loss_{new}(\%)$
1	204.16	186.92	203.78	8.44	0.18
2	204.3	186.92	204.25	8.5	0.02
3	205	186.92	204.95	8.81	0.02
4	208.7	191.93	208.68	8.03	0.009
5	210.8	191.94	210.75	8.94	0.02
6	207.2	191.94	206.74	7.36	0.22
7	208.7	191.96	208.51	8.02	0.09
8	210.57	199.42	210.35	5.29	0.1

Table 5.5: Social welfare comparison in set D of experiments

Experiment	Optimal SW	SW_{over}	SW_{new}	$loss_{over}(\%)$	$loss_{new}(\%)$
1	97.91	87.34	97.91	10.79	0
2	142.96	97.47	142.96	31.82	0
3	142.99	127.14	142.98	11.08	0.006
4	90.3	77.63	90.24	14.03	0.06
5	143.43	137.06	143.42	4.44	0.006
6	143.43	135.57	143.42	5.48	0.006
7	143.43	128.63	143.42	10.31	0.006
8	167.5	127.29	167.46	24.005	0.02

Table 5.6: Social welfare comparison in set E of experiments

Assessment of equilibrium properties in network-wide PSP auction. We carried out a multitude of experiments in order to gain intuition on which of the two strategies is beneficial for bidders as well as whether these strategies constitute equilibria. We ran four types of experiments: a) experiments in which bidders employ exclusively the overbid strategy, b) experiments in which bidders employ exclusively the minimum bid strategy, c) experiments in which one bidder employs the minimum bid strategy and all the others employ the overbid strategy and d) experiments in which one bidder employs the overbid strategy and all the others employ the minimum bid strategy. Note that in each experiment, each bidder employs a specific strategy throughout the auction. We have observed the following:

- We have identified certain cases in which a path bidder benefits by deviating from the overbid strategy and applying the minimum bid strategy, when all the others follow the overbid strategy. In Table 5.7, we present the result of a pair of experiments (of set E) in which bidder 7 obtains a much higher net benefit if he deviates from the overbid strategy to the minimum bid strategy (22.39 versus 2.20). Recall that a and b are the valuation parameters ($\theta_i'(x) = ax + b$). This result is due to the fact that path bidder 7 raises the price under the overbid strategy without excluding other bidders from playing. Thus, all winners pay more under the overbid strategy for almost the same quantity of bandwidth.

- Similarly, we have identified certain cases in which a path bidder benefits by deviating from the minimum bid strategy and applying the overbid strategy, when all the others follow the minimum bid strategy. The result of such a pair of experiments is presented in Table 5.8: bidder 9 obtains a much higher net benefit if he deviates from the minimum bid strategy to the overbid strategy (137.55 versus 80.02). This is due to the fact that bidder 9 raises the price under the overbid strategy and manages to obtain the most of the capacity by excluding the other bidders from playing.

- Considering the previous two remarks, we conclude that none of the two strategies constitutes an equilibrium in the iterated game.

- When all path bidders apply the overbid strategy, the first path bidder to bid, obtains most of the quantity and all the others obtain almost no bandwidth. The underline intuition is the same with that of overbidding in the single-link case. If his valuation is not low, then the path bidder raises the price very much and the others are unable to respond. Note that when the overbid strategy is applied the allocation is affected by the *order* bidders place their bids. An experiment that illustrates this fact is given in Table 5.9. The order of players in this experiment is as follows: 0 1 2 3 11 5 8 9 6 7 10 4, which is then repeated periodically. Note that bidder 11 is the first path bidder to bid and obtains the most of the capacity. The valuation of path bidder 8 is comparable to that of path bidder 11. Nevertheless, since the former places a bid after the latter, he wins no bandwidth at all which is in disagreement with the optimal allocation.

- If all path bidders adopt the minimum bid strategy, then they are all better off than when all apply the overbid strategy, except perhaps for the first path bidder who in certain cases gains almost the entire capacity when all bidders employ the overbid strategy. This exception does not apply if there are single-link bidders with considerable valuation, that bid before the first path bidder. A pair of experiments that illustrates the superiority of the minimum bid strategy is presented in table 5.10. Note that this is a pair of experiments of set C and the order of bidding is as follows: 0 1 2 3 4 5 6 7 8 9 10 11 12 13.

From the above we conclude that it is safer for a path bidder to choose our strategy because, if he is not the first one to play, then he faces the risk of obtaining almost nothing if he did otherwise. In general, overbidding makes sense when one manages to exclude all the other bidders from playing, otherwise he raises prices that will cause him pay ultimately more for a lower quantity. However, the order of players is not determined by the bidders themselves. In the asynchronous implementation of the PSP auction without instant feedback of new bids, a bidder may not be the first one to play but may submit a bid as if he were. Other bidders may act similarly. Only one of them may benefit from this (the first one), but he does not know it prior to bidding. Thus, even if a bidder is willing to risk for being the actual first bidder, he may not benefit from overbidding. He will benefit only if he manages to exclude all other bidders from playing.

Last, we have not considered networks with more than two links because the overbid strategy performs better in the case of two links. Indeed, we have seen that the larger the number of path bidders, the higher the deviation from the optimal social welfare is observed. We expect that the social welfare loss increases as the number of links (and thus the number of path bidders) increases. In fact, we have already seen this in Example 5.3.2. Thus, the negative effect of the

Bidder	**Link**	a	b	**Optimal Alloc**	**Over. Strategy**		**New strategy by bidder 7**	
					Bandw.	NB	Bandw.	NB
0 (seller)	1	0	0.1	0	0	0	0	0
2	1	-0.08	30	4.62	4.24	126.16	4.32	128.44
1 (seller)	2	0	0.1	0	0	0	0	0
3	2	-0.08	6	4.62	4.24	24.31	4.32	24.74
4	1-2	-1	28	0	0	0	0	0
5	1-2	-1	28.2	0	0	0	0	0
6	1-2	-1	28.3	0	0	0	0	0
7	1-2	-1	36	0.37	0.75	2.20	0.67	22.39
8	1-2	-1	26.8	0	0	0	0	0
9	1-2	-1	29.7	0	0	0	0	0
SW				179.07	178.27		178.27	

Table 5.7: A pair of experiments where it is beneficial for bidder 7 to deviate from the overbid strategy

overbid strategy is magnified in the case of large networks. Therefore, it is more interesting to compare the two strategies in the case of two links.

5.6 Concluding Remarks

In this chapter, we discuss certain issues on the PSP auction and propose a new strategy for the network-wide PSP auction. This strategy yields a more efficient outcome, is independent of the bidders' order and is individually beneficial to them, most of the cases. As a future direction, one could investigate various adaptive strategies and how well they perform in the PSP auction. One such strategy could be the following: apply the minimum bid strategy if one observes low demand and the overbid strategy if demand increases. Or alternatively: apply the minimum bid strategy if one realizes that he has a relatively high valuation (compared to his opponents) and otherwise apply the overbid strategy.

Bidder	**Link**	a	b	**Optimal Alloc**	**New Strategy**		**Over. strategy by bidder 7**	
					Bandw.	NB	Bandw.	NB
0 (seller)	1	0	0.1	0	0	0	0	0
2	1	-0.08	6	0	0	0	0	0
1 (seller)	2	0	0.1	0	0	0	0	0
3	2	-0.08	6	0	0	0	0	0
4	1-2	-1	28.3	0.5	0.51	6.40	0	0
5	1-2	-1	28.2	0.4	0.41	4.59	0	0
6	1-2	-1	29.7	1.9	1.88	45.79	0	0
7	1-2	-1	27.5	0	0	0	0	0
8	1-2	-1	26.8	0	0	0	0.03	0.51
9	1-2	-1	30	2.2	2.18	54.72	4.96	80.02
SW				143.43	143.42		137.55	

Table 5.8: A pair of experiments where it is beneficial for bidder 9 to deviate from the minimum bid strategy. Note that bidder 9 is not the first to submit a bid.

Bidder	**Link**	a	b	**Optimal Alloc**	**Over. Strategy**		**New strategy**	
					Bandw.	NB	Bandw.	NB
0 (seller)	1	0	0.1	0	0	0	0	0
2	1	-0.08	6	0	0.25	0.14	0	0
4	1	-1	30	3.2	0	0	3.18	71.04
6	1	-1	11	0	0	0	0	0
1 (seller)	2	0	0.1	0	0	0	0	0
3	2	-0.08	6	0	0.25	0.14	0.007	0.04
5	2	-1	15	3.2	0	0	3.17	24.76
7	2	-1	10	0	0	0	0	0
8	1-2	-1	39	0.4	0	0	0.39	10.71
9	1-2	-1	32	0	0	0	0	0
10	1-2	-1	25	0	0	0	0	0
11	1-2	-1	40	1.4	4.97	132.6	1.42	39.2
SW				204.3	186.92		204.25	

Table 5.9: A pair of experiments to compare bidders' net benefit in the overbid strategy and the minimum bid strategy

Bidder	**Link**	a	b	**Optimal Alloc**	**Over. Strategy**		**New strategy**	
					Bandw.	NB	Bandw.	NB
0 (seller)	1	0	0.1	0	0	0	0	0
2	1	-1	6	0	0	0	0	0
4	1	-1	22	0	0	0	0	0
6	1	-1	19	0	0	0	0	0
8	1	-1	20	0	0	0	0	0
10	1	-1	25	0	0.02	0.15	0	0
12	1	-1	23	0	0	0	0	0
1 (seller)	2	0	0.1	0	0	0	0	0
3	2	-1	6	0	0	0	0	0
5	2	-1	21	0	0	0	0	0
7	2	-1	21	0	0	0	0	0
9	2	-1	21	0	0.02	0.09	0	0
11	1-2	-1	50	2.5	4.98	56.53	2.49	70.40
13	1-2	-1	50	2.5	0	0	2.51	70.98
SW				243.75	237.51		243.74	

Table 5.10: A pair of experiments to compare bidders' net benefit in the overbid strategy and the minimum bid strategy

Chapter 6

Conclusions - Directions for further research

In this dissertation we study auction mechanisms for allocating bandwidth in communication networks. The lack of information about users' demand for bandwidth is the major reason for employing auction mechanisms in communication markets.

First, we formulate and analyze a new hierarchical auction to allocate bandwidth efficiently in a hierarchically structured market. Nowadays, markets involve intermediaries due to the physical and management overheads that arise in direct trading. We consider two levels of hierarchy. The top-level seller allocates the bandwidth to intermediaries. These are the lower-level sellers and allocate their portion of bandwidth to end customers. We take advantage of the distribution of information over all parts involved and coordinate the various trades taking place, so that no one has the incentive to deviate from bidding truthfully. The mechanism provides the efficient overall allocation of bandwidth to the customers as if the top-level seller were to assign this bandwidth directly. We prove that despite the consecutive transactions, customers do not incur extra losses with respect to their net benefits. We also argue that applying efficient mechanisms in each level, this alone does not guarantee overall efficiency. The key issue is demand revelation by the providers. We also define business models in which the intermediate providers are allowed to choose any payment rule to apply in their own local market. We show that the intermediate provider may obtain more profits employing inefficient mechanisms. However, we also prove that if each customer chooses his own provider taking into account the payment rules employed in the various markets, then each provider has the incentive to apply the initial payment rule; for otherwise he would end up with no customers and no profits. Our results are not particular for bandwidth markets; they also apply to other markets for which the hierarchical trading model pertains, such the hierarchical trading of units of other services (e.g. call minutes). The hierarchical auction mechanism can be extended in selling bandwidth through multiple hierarchical levels. Throughout our analysis, we assume that the efficient overall allocation of bandwidth is the ultimate objective of the top-level seller. The study of business models with hierarchical structures in which the top-level seller seeks to maximize his expected profit rather than social welfare is an interesting and challenging direction for further research.

Furthermore we study auctions in which the seller participates as a bidder too. The seller may be a service provider himself and thus he may wish to keep the bandwidth if no one would be willing to pay for it an amount higher than his own valuation. We assess the impact of seller participation on bidding strategies and profits. We derive the equilibrium strategies in first and second-price auctions and compare the expected profits of the seller. An interesting topic of future work is to extend this study to multi-unit auctions, thus giving the owner of the goods the opportunity to benefit by keeping a portion of the bandwidth rather than the entire quantity.

Finally, we revisit the network-wide Progressive Second Price Auction (PSP) proposed by Lazar and Semret, and we introduce a new strategy for path bidders that yields a more efficient outcome than under the overbidding strategy proposed by Lazar and Semret. In fact, the outcome under our strategy is independent of the bidders' order. Moreover, bidders obtain higher expected net benefit by adopting the new strategy in most of the cases. In future research, one could investigate the use of adaptive strategies and how well they perform in the PSP auction. One such strategy could be the following: apply the new strategy if one observes low demand and the overbidding strategy if demand increases. Or alternatively: apply the new strategy if one realizes that he has a relatively high valuation (compared to his opponents) and otherwise apply the overbid strategy.

Another interesting research problem is that of offering and charging contracts for communication services. A service is characterized by two or more parameters such as peak rate, mean rate, etc., thus giving rise to a multidimensional version of the problem of allocating resources. In their selection of service contracts, users should become more sophisticated since they must balance the weights of the different parameters in a way that yields them the maximum net benefit. It also becomes a complicated task for providers to solve the winner determination problem (i.e., derive the allocation resulting in the maximum revenue or the maximum social welfare) as there are more constraints to be taken into consideration on top of the link capacity constraint. The problem of charging such communication services can possibly be solved by resorting to the notion of the effective bandwidth, which provides a measure of resource usage and thus, transforms the multidimensional problem to its one-dimensional counterpart. The study of auction-based charging for services constitutes an interesting yet challenging direction for further research.

The application of auctions as well as of other game theoretic tools in communication markets has attracted much interest recently. This is due to the fact that bandwidth trading can be seen as a game of incomplete information in which players' private information affects strategic behavior. We believe that in the future research work in this area will shed on more problems of strategic interaction among players who compete for resources.

Appendix A

Proof of Proposition 3.2.1

Proof.

Let vector $\vec{x}^h = (x_1^h, \cdots, x_N^h)$ be a concatenation of solutions of the M sub-problems defined in (3.2). For the vector $\vec{x}^* = (x_1^*, \cdots, x_N^*)$ that maximizes the social welfare, we set $\sum_{i \in S_j} x_i^* = q_j'$. Since:

$$\sum_{i=1}^{N} x_i^* = \sum_{j=1}^{M} q_j' = C \text{ and} \tag{A.1}$$

$$\sum_{i=1}^{N} \sum_{k=1}^{x_i^*} \theta_{i,k} = \sum_{j=1}^{M} \sum_{k=1}^{q_j'} v_{j,k}, \tag{A.2}$$

it follows that the vector $\vec{q}'$ is a solution of problem (3.3). Combining this, with definition (3.2), it follows that,

$$\sum_{i \in S_j} \sum_{k=1}^{x_i^*} \theta_{i,k} \leq \sum_{i \in S_j} \sum_{k=1}^{x_i^h} \theta_{i,k}, \tag{A.3}$$

which implies that,

$$\sum_{i=1}^{N} \sum_{k=1}^{x_i^*} \theta_{i,k} \leq \sum_{i=1}^{N} \sum_{k=1}^{x_i^h} \theta_{i,k}. \tag{A.4}$$

Since $\sum_{i=1}^{N} \sum_{k=1}^{x_i^*} \theta_{i,k}$ is the maximum value of the social welfare given by equation (3.1) it follows that,

$$\sum_{i=1}^{N} \sum_{k=1}^{x_i^*} \theta_{i,k} \geq \sum_{i=1}^{N} \sum_{k=1}^{x_i^h} \theta_{i,k}. \tag{A.5}$$

From equations A.4 and A.5 it follows that

$$\sum_{i=1}^{N}\sum_{k=1}^{x_i^*}\theta_{i,k} = \sum_{i=1}^{N}\sum_{k=1}^{x_i^h}\theta_{i,k}. \tag{A.6}$$

Combining this, with the fact the efficient allocation is unique, it follows that the unique solution of the two problems is $\vec{x}^* = \vec{x}^h$, hence they are equivalent.

■

Appendix B

Equilibrium Strategy under Uniform Pricing in the Lower-Level Auction of Example 3.5.2

We consider the complete information game in which each customer knows the marginal valuations of his rivals in his local market and in the other markets too. Each provider reveals his local market's true demand in the top-level auction. We will prove that the following set of strategies constitutes a Nash equilibrium:

- S_1: Customer 1 bids truthfully at each price.
- S_2: Customer 2 bids truthfully up to price 2, reduces demand by 2 units at price 2 and does not change demand up to the termination of the auction.
- S_3: Customer 3 bids truthfully at each price.
- S_4: Customer 4 bids truthfully at each price.

First, assume that customers 1, 3 and 4 bid truthfully at each price. We will prove that customer 2 maximizes his net benefit by adopting strategy S_2. Customer 2 has no incentive to overbid at any price, since he will receive negative profit for the extra units he may obtain and he will pay a higher price for the remaining. (We remind that a bidder pays the final price for each unit he obtains). In addition, customer 2 will not reduce demand up to price 2, since this action does not cause the termination of the auction: rivals' total demand equals 8 which is the capacity. Customer's 2 true demand equals 3 units at price 2. If he reduces demand by 2 units, the auction terminates and he obtains 1 unit as shown in Table 3.14. His net benefit equals $NB_2 = 12 - 2 = 10$. If he reduces demand by 2 units at any price in the interval $(2, 3)$ or by 1 unit at any price in the interval $[3, 5)$, he obtains 1 unit as well but pays a higher price than 2, so his net benefit is less than NB_2. Moreover, he has no incentive to reduce demand by 1 unit at any price in the interval $[2, 3)$ since this action does not cause the termination of the auction. By any other reduction, customer 2 obtains no units at all achieving zero net benefit. If he is truthful at each price, the auction terminates at price 4 and he obtains 2 units as shown in Table 3.12. His net benefit equals then $12 + 5 - 2 * 4 = 9$ which is again less than NB_2. Thus, the optimal strategy for customer 2 provided that customers 1, 3 and 4 bid truthfully, is to reduce demand by two units at price 2 and cause the termination of the auction.

Assume now that customers 2, 3 and 4 bid according to strategies S_2, S_3 and S_4 respectively. Next, we derive the optimal strategy of customer 1. Reasoning as previously, it follows that customer 1 has the incentive neither to overbid at any price, nor to reduce demand up to price 2. At price 2, his rivals' demand equals $1 + 2 + 3 = 6$, thus he bids truthfully for 2 units at this price too. The auction terminates at price 2 and customer 1 receives net benefit $NB_1 = 10 + 4 - 2 * 2 = 10$. Consequently, truthful bidding at each price (i.e., strategy S_1) is the optimal strategy for customer 1.

Similarly, one can show that customers' 3 and 4 optimal strategy is to bid truthfully, so the set $\{S_1, S_2, S_3, S_4\}$ constitutes a Nash equilibrium.

Appendix C

Seller's Bidding Strategy in the First-Price Auction with Seller Participation

Proposition C.0.1 *In a first-price auction with seller participation, truthful bidding is a dominant strategy for the seller provided that $\beta_1(1) = \cdots = \beta_N(1)$, where β_i is bidder's i strategy; that is, provided that bidders' strategies are such that the maximum possible bid of all bidders is the same.*

Proof. Let the seller with valuation x_s submit a bid b_s. Assume further that each bidder i with valuation x_i applies the increasing and differentiable strategy $\beta_i(x_i)$. Let $\beta_1(1) = \cdots = \beta_N(1) \equiv \beta(1)$. Each bidder's valuation is uniformly distributed on [0,1]. We will prove that the seller's optimal strategy β_s^f is to bid truthfully, against any β. We have

$$\Pr[\beta_j(X_j) < b_s] = \begin{cases} \beta_j^{-1}(b_s), & \text{for } b_s \leq \beta(1) \\ 1, & \text{for } b_s > \beta(1) \end{cases} \tag{C.1}$$

The seller's probability of winning is a function of his bid b_s and is given by:

$$\begin{aligned}
P_s(b_s) &= \Pr[\max_{j=1,\ldots,N} \{\beta_j(X_j)\} < b_s] \\
&= \Pr[\beta_j(X_1) < b_s \text{ and } \cdots \text{ and } \beta_j(X_N) < b_s] \\
&= \prod_{j=1}^{N} \Pr[\beta_j(X_j) < b_s] \\
&= \prod_{j=1}^{N} \Pr[X_j < \beta_j^{-1}(b_s)] \Leftrightarrow
\end{aligned} \tag{C.2}$$

$$P_s(b_s) = \begin{cases} \beta_1^{-1}(b_s)\beta_2^{-1}(b_s)\cdots\beta_N^{-1}(b_s), & \text{for } b_s \leq \beta(1) \\ 1, & \text{for } b_s > \beta(1) \end{cases}$$

where we have also used independence of $X_1, \ldots, X_N$. The seller's expected profit E_s equals his valuation x_s if he is the winner; otherwise, it equals the expected value of the maximum bid Y of the N original bids conditional on the seller not being the winner. Therefore,

$$E_s(b_s) = \begin{cases} P_s(b_s) \cdot x_s + [1 - P_s(b_s)] \cdot E[Y|Y > b_s], & \text{for } b_s \leq \beta(1) \\ x_s, & \text{for } b_s > \beta(1) \end{cases} \tag{C.3}$$

where $Y = \max_{j=1,\ldots,N} \{\beta_j(X_j)\}$. Henceforth, to simplify notation, we use $\max_j \{\cdot\}$ instead of $\max_{j=1,\ldots,N} \{\cdot\}$. Assume for now that $b_s \leq \beta(1)$. The distribution function of Y given that $Y > b_s$, is

$$\begin{aligned}
\Pr[\max_j \{\beta_j(X_j)\} \leq y | \max_j \{\beta_j(X_j)\} > b_s] &= \frac{\Pr[b_s \leq \max_j \{\beta_j(X_j)\} \leq y]}{\Pr[\max_j \{\beta_j(X_j)\} > b_s]} \\
&= \frac{\Pr[\max_j \{\beta_j(X_j)\} \leq y] - \Pr[\max_j \{\beta_j(X_j)\} \geq b_s]}{\Pr[\max_j \{\beta_j(X_j)\} > b_s]}.
\end{aligned} \tag{C.4}$$

Reasoning as in the case of (C.2), it follows easily that

$$\Pr[\max_j \{\beta_j(X_j)\} \leq y | \max_j \{\beta_j(X_j)\} > b_s] = \frac{\beta_1^{-1}(y)\beta_2^{-1}(y)\cdots\beta_N^{-1}(y) - \beta_1^{-1}(b_s)\beta_2^{-1}(b_s)\cdots\beta_N^{-1}(b_s)}{1 - \beta_1^{-1}(b_s)\beta_2^{-1}(b_s)\cdots\beta_N^{-1}(b_s)},$$

for any $y \in [b_s, \beta(1)]$.

(C.5)

Differentiating (C.5) with respect to y, we obtain the probability density function f_Y of Y conditional on $Y > b_s$. Thus, for the expected value at large, we have

$$E[Y|Y > b_s] = \int_{b_s}^{\beta(1)} y \cdot f_Y(y)\, dy$$

$$= \int_{b_s}^{\beta(1)} y \cdot \frac{\{[\beta_1^{-1}(y)]'\beta_2^{-1}(y)\cdots\beta_N^{-1}(y) + \cdots + \beta_1^{-1}(y)\beta_2^{-1}(y)\cdots[\beta_N^{-1}(y)]'\}}{1 - \beta_1^{-1}(b_s)\beta_2^{-1}(b_s)\cdots\beta_N^{-1}(b_s)}\, dy$$

$$= \frac{1}{1 - \beta_1^{-1}(b_s)\beta_2^{-1}(b_s)\cdots\beta_N^{-1}(b_s)} \sum_{i=1}^{N} \left\{ \int_{b_s}^{\beta(1)} y[\beta_i^{-1}(y)]' \prod_{j\neq i} \beta_j^{-1}(y)\, dy \right\}. \tag{C.6}$$

Combining (C.2) and (C.3) with (C.6), the seller's expected profit is given by:

$$E_s(b_s) = P_s(b_s)\cdot x_s + \{1 - P_s(b_s)\} \cdot \frac{1}{1 - P_s(b_s)} \sum_{i=1}^{N} \left\{ \int_{b_s}^{\beta(1)} y[\beta_i^{-1}(y)]' \prod_{j\neq i} \beta_j^{-1}(y)\, dy \right\} \tag{C.7}$$

$$= \beta_1^{-1}(b_s)\beta_2^{-1}(b_s)\cdots\beta_N^{-1}(b_s)\cdot x_s + \sum_{i=1}^{N} \left\{ \int_{b_s}^{\beta(1)} y[\beta_i^{-1}(y)]' \prod_{j\neq i} \beta_j^{-1}(y)\, dy \right\}.$$

In order to maximize E_s with respect to b_s, we differentiate the above expression. We have

$$\begin{aligned} E_s'(b_s) &= \sum_{i=1}^{N} [\beta_i^{-1}(b_s)]' \prod_{j\neq i} \beta_j^{-1}(b_s) x_s - \sum_{i=1}^{N} b_s[\beta_i^{-1}(b_s)]' \prod_{j\neq i} \beta_j^{-1}(b_s) \\ &= (x_s - b_s)\left(\sum_{i=1}^{N} b_s[\beta_i^{-1}(b_s)]' \prod_{j\neq i} \beta_j^{-1}(b_s) \right). \end{aligned} \tag{C.8}$$

Therefore,

$$E_s'(b_s) = 0 \Leftrightarrow x_s = b_s, \tag{C.9}$$

while this point corresponds to a maximum because all functions β^{-1} and their derivatives are positive. If indeed $x_s \leq \beta(1)$, then $E_s(b_s)$ as given by (C.3) is maximized for $x_s = b_s$. Notice that the topmost expression in the right-hand side of (C.3) equals x_s for $b_s = \beta(1)$, which implies that if $x_s \leq \beta(1)$, then bidding a quantity $b_s > \beta(1)$ is not beneficial due to (C.9). On the other hand, if $x_s > \beta(1)$, then $E(b_s)$ is increasing in $[0, \beta(1)]$ and constant for $b_s > \beta(1)$. Therefore, any bid $b_s \geq \beta(1)$ is optimal for the seller which includes the case $b_s = x_s$. Therefore, in any case $\beta_s^f(x) = x$. That is, truthful bidding of the seller maximizes his expected profit. ■

Bibliography

[1] Ausubel L. An Efficient Ascending-Bid Auction for Multiple Objects. Revision, 7 August 2002. To appear in American Economic Review.
http://www.ausubel.com/auction-papers/efficient-ascending-auction-r.pdf

[2] Ausubel L. and P. Cramton. Demand Reduction and Inefficiency in Multi-Unit Auctions. Revision, 17 July 2002, University of Maryland.
http://www.ausubel.com/auction-papers/demand-reduction-r.pdf

[3] Binmore K. Fun and Games: A Text on Game Theory. D. C. Heath & Co, 2001, ISBN 0669246034.

[4] Bulow J. and J. Roberts. The Simple Economics of Optimal Auctions. Journal of Political Economy, 97(5): 1060-1090, 1989.

[5] Chakraborty I. and G. Kosmopoulou. Auctions with Shill Bidding. February 2001. To appear in Economic Theory.
http://www.ou.edu/cas/econ/research/2001-9.pdf

[6] Clark D. D. Internet Cost Allocation and Pricing. Internet Economics, L. W. Mcknight and J. P. Bailey, MIT Press, 1997.

[7] Clarke E. Multipart Pricing of Public Goods. Public Choice, 2:19-33, 1971.

[8] Courcoubetis C. and R. Weber. Pricing Communication Networks: Economics, Technology and Modeling. John Wiley & Sons, 2003, ISBN 0470851309.

[9] Courcoubetis C., M. Dramitinos, G.D. Stamoulis. An auction mechanism for bandwidth allocation over paths. In Proceedings of Seventeenth International Teletraffic Congress (ITC),

"Teletraffic Engineering in the Internet Era", Salvador de Bahia, Brazil, 1163 - 1174, December 2-7, 2001.

[10] Courcoubetis C., F. P. Kelly, V. A. Siris, and R. Weber. A Study of Simple Usage-based Charging Schemes for Broadband Networks. Telecommunication Systems, 15(3-4):323-343, 2000.

[11] Groves T. Incentives in Teams. Econometrica, 41:617-631, 1973.

[12] Cramton P. Ascending Auctions. European Economic Review, 42:3-5, 1998.

[13] Gupta A., D. O. Stahl and A. B. Whinston. A Stochastic Equilibrium Model of Internet Pricing. Journal of Economic Dynamics and Control, 21:697-722, 1997.

[14] Kelly F. P. Charging and Accounting for Bursty Connections. Internet Economics, L. W. Mcknight and J. P. Bailey, MIT Press, 1997.

[15] Kelly F. P. On Tariffs, Policing and Admission Control for Multiservice Networks. Operations Research Letters, 15:1-9, 1994.

[16] Krishna V. Auction Theory. Academic Press, 2002, ISBN 012426297X.

[17] Krishna V. and M. Perry. Efficient Mechanism Design. Mimeo, Penn State University, April, 1998.
http://home.uchicago.edu/ preny/papers/Krishna_Perry_1998.pdf

[18] Lalis S., C. Nikolaou, D. Papadakis, and M. Marazakis. Market-Driven Service Allocation in a QoS-Capable Environment. In Proceedings of ICE98, Charleston, South Carolina, USA, 92-100, October 25-28, 1998.

[19] Lazar A. A. and N. Semret. Design and Analysis of the Progressive Second Price Auction for Network Bandwidth Sharing. Revision, 30 September 1999. To appear in Telecommunications Systems, Special Issue on Network Economics.
http://comet.ctr.columbia.edu/ nemo/telecomsys.pdf

[20] Lazar A. A. and N. Semret. The PSP Auction Mechanism for Network Resource Sharing. In Proceedings of Eighth International Symposium on Dynamic Games and Applications, Maastricht, The Netherlands, July 5-8, 1998.

[21] MacKie-Mason J. K. and H. R. Varian. Pricing the Internet. Public Access to the Internet, B. Kahin and J. Keller, MIT Press, 1995.

[22] P. Maillé and B. Tuffin. A Progressive Second Price Mechanism with a Sanction Bid from Excluded Players. In Proceedings of ICQT03, Munich, Germany, 332-341, September 2003.

[23] Maillé P. and B. Tuffin. Multi-Bid Auctions for Bandwidth Allocation in Communication Networks. In Proceedings of IEEE Infocom, Hong-Kong, China, 2004.

[24] Maillé P. and B. Tuffin. Pricing the Internet with Multi-bid Auctions. Technical Report, n. 5235, Irisa, Rennes, France, 2004.
http://www.rennes.enst-bretagne.fr/ pmaille/papiers/maille2004pricingTR.pdf

[25] Myerson R. B. Optimal Auction Design. Mathematics of Operations Research, 6:58-73, 1981.

[26] Reichl P., S. Bessler and B. Stiller. Second-Chance Auctions for Multimedia Session Pricing. In Proceedings of MIPS03, Naples, Italy, November 2003.

[27] Reichl P., B. Stiller, G. Fankhauser. Auction Models for Multi-Provider Internet Connections. In Proceedings of MMB'99, Tier, Germany, September, 1999.

[28] Riley J. G. and W. F. Samuelson. Optimal Auctions. The American Economic Review 71(3):381-392, 1981.

[29] Rothcopf M. H., A. Pekeč and R. M. Harstad. Computationally manageable combinatorial auctions. Management Science, 44(8):1131-1147, 1998.

[30] Sandholm T. Algorithm for Optimal Winner Determination in Combinatorial Auctions. Artificial Intelligence, 135:1-54, 2002.

[31] Semret N. Market Mechanisms for Network Resource Sharing. PhD thesis, Electrical Engineering Department, Columbia University, New York, 1999.

[32] P. B. Tuffin. Revisited Progressive Second Price Auction for Charging Telecommunication Networks. Telecommunications Systems, 20(3):255-263, 2002.

[33] Vickrey W. Counterspeculation, Auctions, and Competitive Sealed Tenders. Journal of Finance, 16:8-37, 1961.

[34] Walsh W. E. and M. P. Wellman. Efficiency and equilibrium in task allocation economies with hierarchical dependencies. In Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence, Stockholm, Sweden, 1999.

[35] Wolfstetter E. Topics in Microeconomics: Industrial Organization, Auctions, and Incentives. Cambridge University Press, 2003.