

Essays on Networks and Market Design

Alexander Teytelboym

Wadham College, University of Oxford

Submitted in partial fulfilment of the requirements of
the degree of D.Phil. in Economics

Trinity Term 2013

Essays on Networks and Market Design

Alexander Teytelboym, Wadham College, University of Oxford

Submitted in partial fulfilment of the requirements of
the degree of D.Phil. in Economics

Trinity Term 2013

Abstract

This thesis comprises four essays in the economics of networks and market design. The common thread in all these essays is the presence of complementarities or externalities. Chapter 2 presents a unified model of networks and matching markets. We build on a contribution by Pycia (2012). We show that strong pairwise alignment of agents' preferences is a necessary and sufficient condition for the existence of strongly stable networks and strongly stable allocations in multilateral matching markets with finite contracts. Strongly stable networks are not necessarily efficient. Although we use a demanding stability concept, strong pairwise alignment allows for complementarities and externalities. In Chapter 3, we generalise the gross substitutes and complements condition introduced by Sun and Yang (2006). Our new condition guarantees the existence of competitive equilibrium in economies with indivisible goods. Competitive equilibrium can be found using an extension of the double-track adjustment process (Sun and Yang, 2009). In this chapter, we also study contract networks (Ostrovsky, 2008). We show that chain-stable contract allocations can exist even in cyclical contractual networks, such as electricity markets, as long as they are appropriately segmented. In Chapter 4, we run a series of experiments to compare the performance of four auctions – first-price, Vickrey, Vickrey-Nearest Rule (Day and Cramton, 2008), and Reference Rule (Erdil and Klemperer, 2010). In our setting, there are two items and three bidders. Two local bidders want an item each, but the global bidder wants both items. We introduce various exposure and package-bidding treatments. We find that the first-price auction always revenue-dominates all the other auctions without any loss in efficiency, strengthening the results of Marszalec (2011). Exposure affects global bidders only in the first-price auction. In other auctions, global bidders often do not take into account the effect of their own bids on their payments. We find no evidence of threshold effects. Finally, in Chapter 5, we develop a new model of online social network formation. In this model, agents belong to many overlapping social groups. We derive analytical solutions for the macroscopic properties of the network, such as the degree distribution. We study the dynamics of homophily – the tendency of individuals to associate with those similar to themselves. We calibrate our model to Facebook data from ten American colleges.

Acknowledgements

I owe an enormous debt of gratitude to my academic supervisor Vincent Crawford whose patience, wisdom, and support has guided me throughout my doctoral work. He first suggested I read David Gale's beautiful book, *The Theory of Linear Economic Models*, which inspired me to work on matching markets. Vince taught me to strive for clarity and rigour in my economic thinking and I am proud to be one of his many students. I would also like to thank Paul Klemperer, who encouraged me to work on auctions and supported me for several years with invaluable comments and advice during seminars. I hope I have not disappointed either of them. I promise to repay them by supervising students the same way they supervised me.

I have been blessed with the opportunity to collaborate with wonderful co-authors. Bassel Tarbush shared his endless energy and enthusiasm for social networks with me, and this became a foundation for a chapter of this thesis. Daniel Marszalec kindly introduced me to the results of his brilliant work on auctions and guided me through the intricacies of experimental design. I hope that some of the work in this thesis will eventually be incorporated into a joint paper.

Elizabeth Baldwin, Dan Beary, Francis Dennig, Jens Gudmundsson, Claudia Herrestahl, Collin Raymond, Andrew Rhodes, as well as Arun Advani, Péter Esó, Fran Flanagan, Meg Meyer, Marek Pycia, John Quah, and Peyton Young provided generous comments and criticisms on different parts of this thesis. Many others have offered feedback at various seminars and conferences. Zaifu Yang kindly suggested an idea that evolved into a chapter of this thesis, and I really benefitted from our discussions in York.

Gillian Coates and Patricia Rice (as well as Vince and Paul) helped me and Daniel obtain enough funding for our experiments. I thank the George Webb Medley Fund,

the OEP Fund, and the John Fell Fund their generous financial support. Hector Solaz was very patient with us for several weeks at CESS. I am also grateful to Julie Minns and Godfrey Keller for everything they have done and continue to do for students in our department.

Bernie Hogan introduced me and Bassel to the analysis of online social networks. He added a new dimension to my research interests. Bob Hahn has given me the confidence to think about policy issues outside the scope of this thesis. Cameron Ott, together with Bassel, selflessly made much of this thesis readable and grammatically sound. I have had many unforgettable dinners with my friends and colleagues at Wadham, Balliol, Nuffield, and University Colleges during which I tried not to talk about my thesis.

My parents gave me and my brother the love for knowledge and have not stopped supporting us for one second. I owe them everything.

Sarah Glatte was with me throughout the last three years. She has offered to read the thesis many times, but I refused. She has always been my harshest critic and my most loving supporter. This thesis is dedicated to her. Now that it is finished, I hope she will read it.

Contents

1	Introduction	15
1.1	Motivation	15
1.2	Overview of the thesis	18
2	Strongly stable networks and matching markets with contracts	25
2.1	Introduction	27
2.1.1	Networks and matching markets	27
2.1.2	Stability	29
2.1.3	Contribution of this chapter	30
2.2	Related literature	32
2.2.1	Networks	33
2.2.2	Matching markets	33
2.3	Networks	36
2.3.1	Ingredients	36
2.3.2	Stability	39
2.3.3	Key assumption about the allocation rule	40
2.3.4	Stability result	40
2.3.5	Efficiency	44
2.4	Multilateral matching markets	46

2.4.1	Ingredients	46
2.4.2	Heterogeneous contracts	49
2.4.3	Stability	51
2.4.4	Main results	52
2.5	Applications of the models	55
2.5.1	Production networks	55
2.5.2	Teams	56
2.5.3	Firm-union bargaining	56
2.6	Conclusion and extensions	57
2.7	Appendix	59
2.7.1	Sufficiency of strong pairwise alignment for a strongly stable contract allocation	59
2.7.2	Necessity of strong pairwise alignment for a strongly stable contract allocation	65
3	Gross substitutes and complements: a generalisation	73
3.1	Introduction	75
3.2	Related literature	77
3.3	Model of a trading economy	78
3.3.1	Ingredients	78
3.3.2	Preferences	79
3.3.3	SC cycles	81
3.3.4	Motivating examples	81
3.3.4.1	Equilibrium with no odd cycles	82
3.3.4.2	No competitive equilibrium with odd cycles	83
3.4	Main result	84
3.5	Application: contractual networks	86

3.5.1	Model	86
3.5.2	Cycles, chains, and stability	89
3.5.3	Result	90
3.6	Discussion	92
3.6.1	Trading networks	92
3.6.2	Importance of quasilinearity	94
3.6.3	Auction	94
3.6.4	Algorithm for testing full substitutability	95
3.7	Conclusion	95
3.8	Appendix	97
4	Auctions for complements: an experimental investigation	101
4.1	Introduction	103
4.2	Related literature	108
4.2.1	Theory	108
4.2.2	Experiments	110
4.3	Theoretical background, experimental design, and hypotheses	112
4.3.1	Model set-up	112
4.3.2	Treatments	112
4.3.3	Auctions	114
4.3.3.1	First-price auction	117
4.3.3.2	Vickrey auction	118
4.3.3.3	Vickrey-Nearest Rule	120
4.3.3.4	Reference Rule	121
4.3.3.5	Illustration of all four auctions	124
4.3.4	Hypotheses	127
4.3.4.1	Statistical tests and methodology	130

4.4	Experimental procedures	130
4.5	Results	132
4.5.1	Baseline results	132
4.5.2	Testing the hypotheses	134
4.5.3	Discussion	146
4.6	Conclusions and extensions	148
4.7	Appendix	150
4.7.1	Subject characteristics and test performance	150
4.7.1.1	Examples of bidding and feedback screens	151
4.7.2	Further tables	152
4.7.2.1	Effect of auction design on bidder surplus and profits	152
4.7.2.2	Profits and revenues in the Combinatorial treatment	154
4.7.2.3	Threshold bidding hypothesis	155
4.7.2.4	Vickrey-Nearest Rule vs. Reference Rule	156
4.7.3	Instructions	158
5	Friending: a model of online social networks	185
5.1	Introduction	187
5.1.1	Homophily	187
5.1.2	Socialising on Facebook	188
5.1.3	Our contribution	189
5.2	Related literature	191
5.3	Model	193
5.3.1	Characteristics of agents	193
5.3.2	Network formation process	195
5.3.3	Interpretation of the model	195
5.3.4	Discussion of the model	197

5.3.5	Relationship to affiliation networks	199
5.4	Theoretical results	200
5.4.1	Degree distribution	201
5.4.2	Assortativity	205
5.4.3	Homophily	206
5.4.3.1	Individual homophily	206
5.4.3.2	Group homophily	206
5.4.3.3	Homophily in our model	207
5.4.3.4	Dynamics of homophily	208
5.5	Simulation results	211
5.6	Data	212
5.7	Tests and empirical observations	213
5.7.1	A representative college	213
5.7.2	All colleges	213
5.8	Model calibration	215
5.8.1	Empirical strategy	215
5.8.2	Results	219
5.9	Discussion	223
5.9.1	Arrival of new nodes	223
5.9.2	Endogenous probability of idleness	225
5.9.3	Preferential attachment	226
5.9.4	Endogenous characteristics	227
5.10	Conclusion	228
5.11	Appendix	230
5.11.1	Simulation algorithm	230
5.11.2	Algorithm for finding robust points in the grid search	231

5.11.3	Data description	232
5.11.4	Results	233
5.11.5	Degree distributions in cleaned and raw data	234
5.11.6	Dynamics of homophily across the grid space	235

List of Figures

2.1	Network represented by a graph Y	28
2.2	Two-sided matching market represented by a graph	28
2.3	Three-sided matching market represented by a graph	29
2.4	Networks represented by hypergraphs Y and Y'	31
3.1	Odd SC cycle	86
3.2	Even SC cycle	86
3.3	Prohibited contract structure	88
3.4	Even cycles in the contractual network	90
3.5	Example of a network that is chain-stable, but not stable	92
4.1	Treatments in the auction experiments	115
4.2	Illustration of all four rules in Case 1 (Example 1)	126
4.3	Illustration of all four rules in Case 5 (Example 2)	128
4.4	Estimated bidding functions for the global bidder in the Exposed and Components treatments of the Vickrey auction	141
4.5	Examples of z-Tree software	151
5.1	Network formation process in the Example	198
5.2	Model as an affiliation network	200
5.3	Relationship between degree and individual homophily indices	210

5.4	Harvard University Facebook network in September 2005	214
5.5	Testing predictions of the model	216
5.6	Best-fitting parameter values by year-group	221
5.7	Illustration of results	222
5.8	Structural properties of the Facebook network at Harvard University	224
5.9	Model with endogenous characteristics as an affiliation network	228
5.10	Degree distributions in cleaned and raw data	234
5.11	Dynamics of homophily	236

List of Tables

1.1	Composition of this thesis	18
2.1	Payoffs in a network that is not strongly stable	43
2.2	Inefficiency in a strongly stable network	45
3.1	Competitive equilibrium with no odd cycles in the SC structure	82
3.2	No competitive equilibrium with an odd cycle in the SC structure	83
3.3	Competitive equilibrium with an odd cycle in the SC structure	97
4.1	Revenue and efficiency	133
4.2	Profits and bidder surplus	133
4.3	Median-difference test results — Revenue	135
4.4	Testing spikes in efficiency	138
4.5	Effect of exposure on global bidders' profit	139
4.6	Global bidders' bids on individual items	140
4.7	Median-difference test results — Global bidders' profit conditional on winning	143
4.8	Combinatorial curse	143
4.9	Frequency of truthful bidding in the Vickrey auction	145
4.10	Subject characteristics	150
4.11	Test results	150

4.12	Median-difference test results — Surplus	152
4.13	Median-difference test results — Global bidders' profit	153
4.14	Median-difference test results — Global bidders' profit	154
4.15	Median-difference test results — Revenue	154
4.16	Mann-Whitney Test — Bid on item X — Local bidder 1	155
4.17	Mann-Whitney Test — Bid on item Y — Local bidder 2	155
4.18	Mann-Whitney Test — Bid on item X — Local bidder 1	156
4.19	Mann-Whitney Test — Bid on item Y — Local bidder 2	157
4.20	Differences in local bidders' bids in the Reference Rule	158
5.1	Degree and class/dorm size	217

Chapter 1

Introduction

1.1 Motivation

In this thesis, we study networks and markets. We are interested in markets that can be “designed” by a social planner. They include auctions (Vickrey, 1961) and matching markets (Gale and Shapley, 1962). We also focus on contractual, trading, and social networks.

Designed markets are used to allocate resources in a variety of settings. For example, every year the U.S. National Resident Matching Program (NRMP) places junior doctors in hospitals (Roth, 1984b). Governments sell spectrum rights to telecommunications service providers through auctions (McMillan, 1994). But when can a successful matching market or auction be designed? In their seminal paper, Kelso and Crawford (1982) showed that if all resources (such as spectrum licences) are substitutes then a stable and efficient allocation of resources is guaranteed to exist and can be found by a well-designed market mechanism. Hatfield and Milgrom (2005) and Hatfield and Kojima (2010) further developed this profound link between matching markets and auctions. They explored matching markets with contracts in which agents specify terms of their relationship in a match. Matching with contracts led to a blossoming in prac-

tical applications of matching markets in the absence of substitutes, such as matching cadets to army units and allocation of airline seat upgrades (Sönmez and Switzer, 2013, Kominers and Sönmez, 2013).

However, in the four chapters of this thesis, we focus on the effects of complementarities and externalities in networks and markets. We say that things are complements if they are valued together more than the sum of their component parts. A pair of shoes is worth something to us, but either shoe is worth nothing. In his review of John Hicks and Roy Allen’s work, Paul Samuelson observed

The time is ripe for a fresh modern look at the concept of complementarity . . . The last word has not yet been said on this ancient preoccupation of literary and mathematical economists. The simplest things are often the most complicated to understand fully (Samuelson, 1974, p. 1255).

Complementarities create theoretical and practical challenges in market design. Kelso and Crawford (1982) showed that stable matchings may fail to exist when complementarities are present. Various other theorists pointed out difficulties in designing matching markets with complementarities and externalities (Sasaki and Toda, 1996, Echenique and Yenmez, 2007). Yet theoretical challenges did not make the problem disappear in practice: since the early 1980s, couples could apply for residence in nearby hospitals through the NRMP mechanism. This introduced complementarities in the mechanism and some difficulties in the implementation of stable allocations (Roth and Peranson, 1999).

Moreover, during the U.S. spectrum auctions in the 1990s, theorists began to pay attention to the practical effects of complements in auction design (Krishna and Rosenthal, 1996, Levin, 1997). The heterogeneity of licences, technologies, and locations meant that different bidders regarded different sets of licences as complements or substitutes. However, the presence of complementarities in standard auctions, such as first-price and Vickrey, was shown to cause inefficiencies and result in low revenues.

Auction theorists filled this gap by proposing auctions that could accommodate complementarities and many of their designs are already used in spectrum auctions around the world (Ausubel and Milgrom, 2002, Day and Milgrom, 2008, Day and Cramton, 2008, 2012, Erdil and Klemperer, 2010).

These developments in matching markets and auctions coincided with the integration of network theory into economics (Jackson and Wolinsky, 1996). Network economics is devoted to understanding how the structure of our connections affects economic and social outcomes. Externalities, such as peer effects, are central to theoretical and empirical analyses in network economics. Like matching theory, network economics has a social design perspective. Its goal is to understand how surplus can be allocated on a network to ensure stable and socially optimal outcomes.

However, many real-world networks are not designed. Nevertheless, their properties affect outcomes of network processes. For example, various individual behaviours, such as obesity and voting, spread through social networks (Christakis and Fowler, 2007, Bond et al., 2012). Therefore, another focus of network economics has become the study of how social and economic networks form and evolve (Jackson and Rogers, 2007).

In recent years, the economic analysis of networks and market design has been merging. Ostrovsky (2008) developed a theoretical model of matching in production networks with complements. Baccara et al. (2012) empirically studied how network externalities affect outcomes in assignment markets. Intersecting computer science and economics, Askalidis et al. (2013) looked at whether it is possible to improve the efficiency of matching markets for hospitals and doctors using the information about the underlying social network. However, many issues relating to the presence of complementarities and externalities in networks and market design remain unresolved. This thesis makes a modest attempt to fill this gap.

1.2 Overview of the thesis

Table 1.1 provides an overview of topics covered in this thesis.

Table 1.1: Composition of this thesis

Chapter	Networks	Markets	Complementarities	Externalities
2	★	★	★	★
3	★	★	★	
4		★	★	
5	★			★

The main part of the thesis begins with two theory chapters. Chapter 2 presents a unified model of networks and matching markets with contracts. We build on a recent contribution by Pycia (2012). In his model, Pycia finds a necessary and sufficient condition for the existence of a (unique) stable allocation in many-to-one matching markets. This condition is called *pairwise alignment* and states that for any pair of matches, the preferences of agents who belong to these matches should be aligned. We extend his model in two directions using a new condition called *strong pairwise alignment*. First, we consider a network model in which the social planner can enforce a rule that allocates surplus. We then show that strong pairwise alignment is necessary and sufficient for the existence of strongly stable networks (Dutta and Mutuswami, 1997). Second, we consider a many-to-many multilateral matching market with finite contracts. In this model, agents have preferences over contracts with agents who belong to other sides of the market. We show that even in this very general model strong pairwise alignment is also necessary and sufficient for the existence of strongly stable contract allocations (Hatfield and Kominers, 2011a). Strong stability in both models is a demanding stability concept that requires robustness to arbitrary deviations of groups of agents. However, strong pairwise alignment allows for the presence of complementarities

and externalities in our models. The model has a variety of potential applications including production networks, firm-union bargaining, and teams within organisations.

Chapter 3 looks at the existence of competitive equilibrium in a market for indivisible goods and of stable contractual allocations in production networks. First, we generalise *gross substitutes and complements* preferences (Sun and Yang, 2006). In Sun and Yang’s framework, there are two sets of substitutable goods, e.g. shirts and jackets, and agents regard any shirt and jacket as complements. In our model, there are many disjoint sets of substitutable goods and agents view goods from any two sets as complements. We show that whenever these complementary preferences do not form any *odd cycles*, competitive equilibrium is guaranteed to exist. An odd cycle is present when one person views trousers and shirts as complementary, another – shirts and jackets, and yet another – jackets and trousers. We show that the double-track auction developed by Sun and Yang (2009) can be extended to find the competitive equilibrium allocation and prices. Second, we consider contract networks in which firms have *same-side substitutable* and *cross-side complementary* preferences over contracts (Ostrovsky, 2008). Firms regard all their downstream/upstream contracts as substitutes, but any downstream and upstream contracts as complements. Unlike Chapter 2, we focus on the existence of contract allocations that satisfy a rather weak stability concept called *chain* (or pairwise) *stability*. We prove that whenever production networks are appropriately segmented and there are no odd contractual cycles, chain-stable contract allocations always exist.

In Chapter 4, we turn to an experimental study of auctions for complements. We study a simple setting with two items and three bidders. One of the bidders (global) is interested in buying both items, while the other two bidders (local) are only interested in the individual items. We consider four auctions: first-price and Vickrey as well as two minimum-revenue core-selecting (MRCS) auctions called the Vickrey-Nearest Rule

(Day and Cramton, 2008) and the Reference Rule (Erdil and Klemperer, 2010). If the items were substitutable, the auctioneer would face no serious design difficulties because the Vickrey auction is efficient and stable in this case. However, the Vickrey auction may raise low revenues and the first-price auction is often inefficient when complementary items are on sale. MRCS auctions are an attempt to overcome these problems. Marszalec (2011) showed that the Vickrey auction performs poorly in terms of revenue *and* efficiency in this setting compared to the other auctions. We look at four treatments in each auction, testing whether auction performance is affected by *exposure* of the global bidder and opportunities to place package bids. The global bidder is exposed whenever he risks ending up with just one of the two complementary items he is interested in. We find that the first-price auction revenue-dominates all the other auctions without much loss in efficiency in every treatment. Global bidders are only affected by exposure in the first-price auctions. However, in other auctions, global bidders often fail to take into account the effect of their own bids on their payments. Local bidders bid very close to their value in the Vickrey auction as predicted by theory. However, we find no evidence of their attempting to free-ride on each other's bids when the global bidder is not exposed (this is known as the *threshold* effect) in any auction.

Finally, Chapter 5 is devoted to a theoretical and empirical study of online social networks. We look at how agents' characteristics can affect the network formation process. In our model, agents are endowed with various social categories, such as gender or a hobby. Agents randomly form their links within particular social groups. Agents who share more social groups are more likely to meet. Since these social groups overlap, the resulting friendship pattern is complex. These social groups also create an interesting externality – agents in smaller social groups create friendships at a slower rate over time. We find analytical solutions to various properties of the resulting social network, such as the degree distribution. In particular, we show how *homophily* –

the tendency of similar agents to be linked – evolves over time. We then calibrate our model to data from the online social network *Facebook* that covers ten U.S. universities. Although our model is very parsimonious, it fits the data well. We discuss various extensions to our model and suggest how it could be used by a social planner.

References

- Askalidis, G., N. Immorlica, A. Kwanashie, D. F. Manlove, and E. Pountourakis (2013). Socially stable matchings. Technical Report arXiv:1302.3309 [cs.GT].
- Ausubel, L. M. and P. R. Milgrom (2002). Ascending auctions with package bidding. *Frontiers of Theoretical Economics* 1(1), 1–42.
- Baccara, M., A. İmrohoroğlu, A. J. Wilson, and L. Yariv (2012). A field study on matching with network externalities. *American Economic Review* 102(5), 1773–1804.
- Bond, R. M., C. J. Fariss, J. J. Jones, A. D. I. Kramer, C. Marlow, J. E. Settle, and J. H. Fowler (2012, September). A 61-million-person experiment in social influence and political mobilization. *Nature* 489, 295–298.
- Christakis, N. A. and J. H. Fowler (2007). The spread of obesity in a large social network over 32 years. *New England Journal of Medicine* 357(4), 370–379.
- Day, R. W. and P. C. Cramton (2008). Quadratic core-selecting payment rules for combinatorial auctions. Working paper, University of Maryland.
- Day, R. W. and P. C. Cramton (2012). Quadratic core-selecting payment rules for combinatorial auctions. *Operations Research* 60(3), 588–603.
- Day, R. W. and P. Milgrom (2008). Core-selecting package auctions. *International Journal of Game Theory* 36(3-4), 393–407.
- Dutta, B. and S. Mutuswami (1997). Stable networks. *Journal of Economic Theory* 76(2), 322–344.
- Echenique, F. and M. B. Yenmez (2007). A solution to matching with preferences over colleagues. *Games and Economic Behavior* 59(1), 46–71.
- Erdil, A. and P. Klemperer (2010). A new payment rule for core-selecting package auctions. *Journal of the European Economic Association* 8(2-3), 537–547.
- Gale, D. and L. S. Shapley (1962). College admissions and the stability of marriage. *American Mathematical Monthly* 69(1), 9–15.
- Hatfield, J. W. and F. Kojima (2010). Substitutes and stability for matching with contracts. *Journal of Economic Theory* 145(5), 1704–1723.
- Hatfield, J. W. and S. D. Kominers (2011a). Contract design and stability in matching markets. Working paper, Harvard Business School.

- Hatfield, J. W. and P. Milgrom (2005). Matching with contracts. *American Economic Review* 95(4), 913–935.
- Jackson, M. O. and B. W. Rogers (2007). Meeting strangers and friends of friends: How random are social networks? *American Economic Review* 70(3), 890–915.
- Jackson, M. O. and A. Wolinsky (1996). A strategic model of social and economic networks. *Journal of Economic Theory* 71(1), 44–74.
- Kelso, A. S. and V. P. Crawford (1982). Job matching, coalition formation, and gross substitutes. *Econometrica* 50(6), 1483–1504.
- Kominers, S. D. and T. Sönmez (2013). Designing for diversity in matching. Working paper, Boston College.
- Krishna, V. and R. Rosenthal (1996). Simultaneous auctions with synergies. *Games and Economic Behavior* 17(1), 1–31.
- Levin, J. (1997). An optimal auction for complements. *Games and Economic Behavior* 18(2), 176–192.
- Marszalec, D. (2011). *Essays on Auctions*. Ph. D. thesis, Nuffield College, University of Oxford.
- McMillan, J. (1994). Selling spectrum rights. *Journal of Economic Perspectives* 8(3), 145–162.
- Ostrovsky, M. (2008). Stability in supply chain networks. *American Economic Review* 98(3), 897–923.
- Pycia, M. (2012). Stability and preference alignment in matching and coalition formation. *Econometrica* 80(1), 323–362.
- Roth, A. E. (1984b). Stability and polarization of interests in job matching. *Econometrica* 52(1), 47–58.
- Roth, A. E. and E. Peranson (1999). The redesign of the matching market for American physicians: Some engineering aspects of economic design. *American Economic Review* 89(4), 748–780.
- Samuelson, P. (1974). Complementarity: An essay on the 40th anniversary of the Hicks-Allen revolution in demand theory. *Journal of Economic Literature* 12(4), 1255–1289.
- Sasaki, H. and M. Toda (1996). Two-sided matching problems with externalities. *Journal of Economic Theory* 70(1), 93–108.

- Sönmez, T. and T. Switzer (2013). Matching with (branch-of-choice) contracts at the United States Military Academy. *Econometrica* 81(2), 451–488.
- Sun, N. and Z. Yang (2006). Equilibria and indivisibilities: Gross substitutes and complements. *Econometrica* 74(5), 1385–1402.
- Sun, N. and Z. Yang (2009). A double-track adjustment process for discrete markets with substitutes and complements. *Econometrica* 77(3), 933–952.
- Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* 16(1), 8–37.

Chapter 2

Strongly stable networks and matching markets with contracts

Abstract

We present a unified model of contractual networks and matching markets with finite contracts. We offer a necessary and sufficient condition for the existence of strongly stable networks (Dutta and Mutuswami, 1997) and strongly stable contract allocations in multilateral matching markets. Strongly stable allocations are robust to profitable deviations by groups of agents. Our condition allows for the presence of externalities across and complementarities within contracts. Substitutable preferences are not necessary for the existence of strongly stable allocations. However, following Pycia (2012), we require substantial alignment of preferences between sets of contracts in order to guarantee strong stability.

Keywords: networks, externalities, multilateral matching, stability, complementarities, contracts.

JEL Classification: *C7, D4, D7, L14*

2.1 Introduction

Matching markets and networks are omnipresent in economic activity.¹ A network is a structure of connections. For example, a network may represent how firms collaborate on research and development (R&D) projects. A multilateral matching market consists of multiple types of heterogeneous agents that have preferences over productive matches with agents of the other types. Canonical matching markets – for hospitals and medical interns – are two-sided. Although matching markets and networks share a common mathematical structure, they are usually explored separately in the economics literature. This chapter offers a unified theoretical approach which brings new insights in both fields.

2.1.1 Networks and matching markets

We can use graphs to describe both networks and outcomes in matching markets. A vertex (or a node) of a graph represents an agent, such as a firm, and an edge (or a link) between two vertices represents a resulting productive connection or a match. The links encode the nature of the relationships between the nodes, such as an employment contract, an R&D venture, or an allocation of tasks in a team. Figure 2.1 represents a network Y comprising ten agents and two connected *components*, C and C^* , i.e. subsets of nodes in which any two nodes are connected to each other by a path. Figures 2.2 and 2.3 represent outcomes in two- and three-sided matching markets respectively. Here, agents belong to different sides of the market e.g. venture capital funds (f), inventors (g), and universities (h).

The most important similarity between graphs representing networks and matching markets is that all links are formed cooperatively and any agent can break a link without the consent of affected agents. Moreover, all links and matches are indivisible: a doctor

¹Roth and Sotomayor (1990) offer an excellent treatment of the classic results in matching theory. Jackson (2008) and Goyal (2009) present recent surveys of the networks literature.

Figure 2.1: Network represented by a graph Y






Figure 2.2: Two-sided matching market represented by a graph



cannot be matched to half of a hospital in the same way that firms cannot sign half a contract.

There are, however, some important differences. First, in networks, any group of agents can form links among themselves (Figure 2.1). In matching markets, however, agents usually form productive links only with agents on the other sides of the market. Thus any group of firms can in principle collaborate on R&D, but a doctor would only want to be matched to a hospital. Graphs can also represent the type of market we are interested in. For example, if there are two types of agents and if agents of one type g are allowed to form only one link with an agent on the other side of the market, but agents in f can form many links with agents in g , then the resulting bipartite graph would represent an outcome of any two-sided many-to-one market (Figure 2.2 gives an example of a matching in this type of (g -to- f) matching market). A multilateral market in which agents can form any number of bilateral links is represented in Figure 2.3.

Figure 2.3: Three-sided matching market represented by a graph



A second difference is that most of the matching literature assumes that matches do not create any externalities affecting agents not involved in the match.² For instance, it is common to assume that profits of a firm are not directly affected by the workers other firms hire. In the networks literature, however, agents' payoffs depend on the value of the component to which they belong. The value of each component of network Y in Figure 2.1 and the payoffs of its agents may depend on its structure, i.e. the configuration of links between agents in that component. However, agents do not necessarily need the consent of all other agents in the component to alter its structure. Hence, agents' payoffs may be affected by the links of others. For example, when nodes represent firms, links represent R&D ventures, and a component represents an industry, we can naturally interpret the effect that new links in the component have on the payoffs of all its agents as the existence of technological spillovers. The network structure of the ventures in the industry determines its total surplus and hence how the spillovers and externalities affect the profits of all its firms.

2.1.2 Stability

In a fully decentralised economy, networks and matchings would form according to agents' private incentives. However, it is natural to consider whether a social designer can improve upon decentralised outcomes by coordinating the matching market or re-

²There are notable exceptions, such as Sasaki and Toda (1996) and Bando (2012).

allocating the surplus on the network. In order to sustain a socially desired outcome, the social designer must consider which matchings and networks are stable, i.e. immune to deviation by agents (Gale and Shapley, 1962, Jackson and Wolinsky, 1996). For example, centralised matching markets in which outcomes are not stable tend to unravel and break down over time as agents seek profitable opportunities outside the prescribed allocation (Roth, 1991). Similarly, if the social designer has the power to redistribute the value (profit) of a network according to an allocation rule, he must consider whether the resulting network will be stable.


This chapter focuses on the existence of networks and allocations in matching markets that satisfy a demanding stability condition. This is particularly important when the market is complex, communication is easy, or there are numerous productive opportunities. Networks and matching market outcomes that satisfy this stability condition are particularly unlikely to unravel, which make them appealing to a social designer.

2.1.3 Contribution of this chapter

The present chapter builds on the seminal results of Pycia (2012). He proves that, if for any two feasible matches m and m' , all agents who belong to *both* matches have the same preferences over m and m' (a condition called *pairwise alignment*), then the many-to-one matching market (Figure 2.2) has a stable allocation.

The contribution of this chapter is twofold. First, we examine rules that stipulate how agents allocate divisible surplus on networks. Our networks are represented by hypergraphs – graphs in which there may be links between more than two nodes. In Figure 2.4, we illustrate two hypergraphs Y and Y' in which multilateral links are represented by dashed lines. We present a necessary and sufficient condition for the existence of *strongly stable* networks i.e. those that are immune to deviations by groups of agents who may want to keep some of their existing links, to form new links only among themselves, and obtain strictly higher payoff (Dutta and Mutuswami, 1997). This condition,

Figure 2.4: Networks represented by hypergraphs Y and Y'



which we call *strong pairwise alignment*, states that agents on the intersection of two components (from two different networks) must have the same preferences over these components. Let us illustrate strong pairwise alignment. In Figure 2.4, agents 1, 2, and 3 are on the intersection of components C and C' . Strong pairwise alignment implies that agent 1 gets higher payoff in component C than in C' if and only if agents 2 and 3 also do. In particular, equal split of component value (the egalitarian solution) satisfies strong pairwise alignment. In order to see this, note that if value is split equally in every component, then agents would strictly prefer to be members of a different component if and only if it has higher per capita value. However, while more general Nash-bargaining allocation rules also satisfy strong pairwise alignment, the celebrated Myerson (1977) value does not.

We then extend Pycia's many-to-one setting without contracts to a multilateral matching market with finite contracts (Hatfield and Kominers, 2011b). In a multilateral matching market, there are many types of agents who can form (multilateral) links with agents of the other types. Agents can form several contracts with the same or different sets of agents. Thus a multilateral market with contracts is even more general than what we presented in Figures 2.2 or 2.3, in which links represent contracts that

specify how a relationship among agents is organised. A finite set of feasible contracts implies that surplus is neither fully divisible nor transferrable between agents involved in a match. We also give a natural interpretation of the “rich preference domain” conditions introduced by Pycia (2012) by demanding that contracts are sufficiently heterogeneous. Contracts express many dimensions of a relationship between agents. For example, in a contract between a venture capital fund and an inventor, the contract could state what the inventor’s equity, debt, and voting rights are. We argue that if contracts are sufficiently heterogeneous then agents’ preference domain over these contracts is likely to be rich. If contracts contain different terms, then the ranking of contracts may vary substantially across preference profiles. In this case, strong pairwise alignment is also necessary and sufficient for the existence of strongly stable contract allocations in multilateral matching markets in all preference profiles. As in the networks model, strong pairwise alignment allows for complementarities within contracts and externalities across contracts. This puts in perspective some recent results in matching theory and shows that, in multilateral matching markets, substitutable preferences are not necessary for the existence of strongly stable contract allocations.

The rest of this chapter is organised as follows. Section 2.2 gives a brief overview of related work. Section 2.3 presents the networks model. Section 2.4 extends this to a multilateral matching market model with finite contracts. Section 2.5 provides possible applications of the model. Section 2.6 concludes. Proofs are in the Appendix.

2.2 Related literature

We now relate the chapter to the existing results in the networks and matching markets literatures.

2.2.1 Networks

In their influential paper, Jackson and Wolinsky (1996) considered pairwise stable networks i.e. those which are immune to deviations by pairs of agents. However, even if a network is pairwise stable, it may not be stable with respect to deviations by groups of agents. Dutta and Mutuswami (1997) were the first to address this concern by analysing strongly stable networks. These networks, when represented as strategic link formation games, are supported by a Strong Nash equilibrium.³ Dutta and Mutuswami (1997) show that it may not be possible to construct efficient (surplus-maximising) and strongly stable networks with pre-arranged allocation rules. We give necessary and sufficient conditions for the existence of strongly stable networks and show that, in general, the conflict between stability and efficiency persists in our model. Jackson and van den Nouweland (2005) also consider strongly stable networks. They find that the egalitarian allocation rule is a necessary and sufficient condition for the existence of strongly stable networks, but their definition of strong stability is even stronger than the one adopted here and by Dutta and Mutuswami (1997).

2.2.2 Matching markets

Matching markets consist of agents who belong to different sides of the market. Agents want to form productive links with agents on other sides of the market. For example, in a two-sided matching market, firms match with workers by offering them a wage. Most analytical approaches impose restrictions on the agents' (strict) individual preferences over the other side of the market, which guarantee a stable matching (Gale and Shapley, 1962).⁴ Substitutability, introduced by Kelso and Crawford (1982),

³In a non-cooperative game, a Nash equilibrium is “strong” if no coalition of agents can cooperatively deviate to benefit all its members while taking the actions of agents outside the coalition as given (Aumann, 1959).

⁴*Additive separability* (or *responsiveness*) of preferences guarantees the existence of stable allocations in one-to-one and many-to-one matching markets (Gale and Shapley, 1962, Crawford and Knoer, 1981).

has been the most influential restriction on preferences because of its natural economic interpretation. Under substitutability, we assume that if a matching option is accepted, then this option is also accepted from a smaller choice set. In the context of many-to-one matching markets, Kelso and Crawford (1982) showed that substitutability guarantees stable allocations (equivalent to the core), which can be found using their salary-adjustment process. Blair (1988) showed a similar result for many-to-many matching markets.

More recently, several papers have generalised the insights of Gale and Shapley (1962) and Kelso and Crawford (1982). In particular, Hatfield and Milgrom (2005) reintroduced the idea of a contract between workers and firms, described originally by Crawford and Knoer (1981). A contract describes various dimensions of a productive relationship e.g. wage, benefits, sick pay, and so on. Hatfield and Milgrom (2005) showed that if all firms regard all contracts with workers as substitutes then core allocations exist in a many-to-one market. This, in turn, led Hatfield and Kojima (2010) to observe that core allocations may exist under even weaker conditions.⁵

Whereas in one-to-one matching markets, one typically considers pairwise stable outcomes (cf. Jackson and Wolinsky, 1996), more complex matching markets give rise to stronger stability concepts.⁶ For example, core allocations in many-to-many matching markets are immune to deviations of agents who drop all their previous matches and form new ones *only* among themselves. Recently, Hatfield and Kominers (2011a) proved that substitutability is effectively necessary for certain stronger stability concepts (which are described in detail in footnote 18) in many-to-many two-sided matching markets with contracts. In this chapter, we use the strongest stability concept found

⁵In fact, under the substitutability assumption, matching markets with contracts are no more general than matching markets with wages (Echenique, 2012, Kominers, 2012). Under the assumptions in this chapter, this is not true (Flanagan, 2013).

⁶Jackson and Watts (2010) explored an interesting setting in which agents first match and then engage in a non-cooperative game with their matched partners.

in the matching markets literature.

In many economic contexts, however, substitutability is a rather strong assumption on individual preferences. Suppose a venture capital fund f faces two inventors g_1 and g_2 and ideally would fund both. Now suppose inventor g_2 goes bankrupt and is no longer available. Substitutability implies that f still wants to fund g_1 , whereas f 's preferences could be such that she wants to fund neither. This could be because f regards inventors g_1 and g_2 as complementary: the technologies they developed are only marketable when sold together. As Kelso and Crawford (1982) showed, without substitutability, stable outcomes may not exist. Nevertheless, complementarities may be introduced into matching markets while preserving stable allocations. Sun and Yang (2006) showed that if firms regard workers as substitutes and machines as substitutes but a worker and a machine as complements (what they called *gross substitutes and complements*), then core allocations will exist in the market.⁷ Externalities, such as peer effects (preferences over matchings with agents on one's own side of the market), also upset substitutability and usually preclude the existence of stable matchings.⁸

Many matching markets are multi-sided. For example, a typical R&D venture may include a fund, a group of inventors, and a university (Macho-Stadler et al., 2008). However, the existence of stable matchings in multilateral matching markets is still an open problem. Alkan (1988) showed that whenever agents want to be matched in threesomes and have preferences over two types of agents, stable matchings may not exist. While Alkan's result holds even for additively separable preferences, Hatfield and Kominers (2011b) showed that stable matchings can be found if continuous prices are introduced in the model. Our model is more general than Alkan's, and we give necessary

⁷Ostrovsky (2008) and Hatfield and Kominers (2012) applied the insight of gross substitutes and complements to models of trading networks. Baldwin and Klemperer (2013) showed that competitive equilibrium may exist in markets for indivisible goods with substitutes and complements. Chapter 3 of this thesis extends these ideas further.

⁸In many-to-one matching markets, algorithms have been developed to find a core allocation – whenever it exists – in the presence of peer effects (Echenique and Yenmez, 2007, Kominers, 2010).

and sufficient conditions for the existence of strongly stable contract allocations in the absence of continuous prices.

Several papers prior to the contribution of Pycia (2012) showed that imposing restrictions on the agents' preference profile, rather than on individual preferences, can accommodate complementarities and externalities while preserving the existence of stable outcomes and avoiding substitutability. These conditions typically reflect a degree of similarity of agents' preferences. Dutta and Massó (1997) and Revilla (2007) showed that if workers care about their co-workers, but their preferences over firms dictate their overall preferences between firms and co-workers (i.e. preferences are lexicographic), then the many-to-one matching market has a core allocation. Dimitrov and Lazarova (2011) found that lexicographic preferences are necessary and sufficient for a non-empty core in a related two-sided coalitional matching problem.

2.3 Networks

We begin by characterising strongly stable networks.

2.3.1 Ingredients

N is a finite set of agents, where $|N| = n$. A network Y (mathematically represented by a *hypergraph*, see Figure 2.4) on N is a family of subsets of N , which are at least of size two. Denote by $Y^N \equiv \{y \in 2^N \mid |y| \geq 2\}$ the set of *all* subsets of N that have size at least two. Any $y \in Y$ is called a *hyperedge*.⁹ For any $S \subseteq N$, define Y^S analogously. The interpretation of the model is straightforward. Any y represents a *contract* between the firms in y . Any two firms i and j that belong to contract y can be thought of as partners. Y^N is the set of all possible contracts between agents in N . All connections between firms are voluntary and reciprocal. We use hypergraphs in order to capture

⁹Hypergraphs generalise graphs in which all hyperedges $y \in Y$ have cardinality 2 and represent edges connecting vertices. All our results apply to graphs.

the fact that some of the agents already linked by a contractual commitment may want to form separate contracts (Myerson, 1980, Slikker et al., 2000, Kongo, 2011). For example, three car manufacturers may be cooperating on engine development and two of them may work on a separate project on transmission. Therefore, every contract is defined by the set of firms which are part of it. This assumption will be relaxed in the next section by allowing the same firms to sign multiple contracts. A component of the network may represent an industry or a sector. Denote the contracts to which agent i belongs (i 's neighbourhood) as $Y_i \equiv \{y \in Y | i \in y\}$.

The set of all possible hypergraphs on N is $\mathcal{Y} = \{Y | Y \subseteq Y^N\}$. Any two agents i and j are *connected* in a hypergraph Y if there exist finite sequences of agents $(i, i_1 \dots j)$ and of hyperedges $(y_0 \dots y_K)$, such that $i = i_0$ and $i_K = j$, and for any $k = 1, \dots, K$, $\{i_{k-1}, i_k\} \subseteq y_k$. If any two distinct agents are connected, we say the hypergraph Y is connected. If the hypergraph Y is not connected, then N can be partitioned into connected *components*. This partition is defined as $\Pi(Y) \equiv \{\{i\} \cup \{i \text{ is connected to } j \text{ in } Y\} | i \in N\}$. For any $S \in \Pi(Y)$, let us define a subhypergraph of $Y(S) \equiv \{y \in Y | y \subseteq S\} \equiv Y \cap Y^S$. The set of (non-empty) components on Y is therefore $C(Y) \equiv \{Y(S) | S \in \Pi(Y), |S| \geq 2\}$.

The *value function* $v : \mathcal{Y} \rightarrow [0, \infty)$ defines the value for each network. Without any loss of generality, we assume that $v(\emptyset) = 0$. We also assume that the value function is *component-additive*: $v(Y) = \sum_{C \in C(Y)} v(C)$ for all $Y \in \mathcal{Y}$. The set of all component-additive value functions is \mathcal{V} . Since the value function is component-additive, it allows externalities within components – a firm may benefit from the R&D activities done by firms within its industry – but not across components. The value function is not assumed to be anonymous, therefore the value of the network may depend on factors other than the network topology. This is natural in many economic contexts. Consider a star network in which the center node is the government and the peripheral nodes are

defense companies. Defense companies do not cooperate with one another. Effectively, the government is awarding various contracts to competing defense companies. Let us switch the central position of the government and the peripheral position of a defense company. Now, the government has one monopolistic supplier which contracts with the other peripheral firms downstream. The total surplus in these two networks will be different even though the network topology is unchanged.

For every network and every value function, the agents' payoffs are determined by an *allocation rule* $\gamma : \mathcal{Y} \times \mathcal{V} \rightarrow \mathbb{R}^n$ fixed by the social designer (Dutta and Mutuswami, 1997, Jackson, 2005, Jackson and van den Nouweland, 2005). The set of all allocation rules is Γ .¹⁰ We make the following assumptions about the allocation rule:

1. *Regular* – γ is strictly increasing and continuous. Define $q : \mathcal{Y} \times [0, \infty) \rightarrow \mathbb{R}^n$, such that $q(Y, v(C)) = \gamma(Y, v(C))$ for all $Y \in \mathcal{Y}$, $v \in \mathcal{V}$, and $C \in C(Y)$. We assume that $\lim_{w \rightarrow \infty} q_i(Y, w) = \infty$ for all $w = v(C)$ and $i \in C$. This assumption states that as the value of the network component goes to infinity, so do the payoffs of all agents.
2. *Anonymous* – for any $v \in \mathcal{V}$, $Y \in \mathcal{Y}$ and permutation π on N , $\gamma_{\pi(i)}(Y^\pi, v^\pi) = \gamma_i(Y, v)$, where the $v^\pi(Y) = v(Y^{\pi^{-1}})$ for each $Y \in \mathcal{Y}$. In general, anonymous allocation rules do not imply that agents in completely symmetric networks get the same payoffs. Agents' payoffs may depend on their characteristics as well as their position in the network as in the example of the government and defense companies above (Jackson, 2005).

¹⁰As defined, Γ is very general. In particular, some allocation rules in Γ do not rule out the possibility that an allocation in a particular hypergraph depends on the value of other hypergraphs. These allocation rules may seem economically implausible. Two interesting classes of allocation rules are contained in our definition. In the first class, we only allow the allocation rule to depend on the value of the hypergraph; hence the rule is defined by a function composition $\gamma : \mathcal{Y} \rightarrow \mathcal{V} \rightarrow \mathbb{R}^n$. In the second class, we impose the condition that for any two value functions which give a particular hypergraph the same value the allocation rule must be the same. Specifically, for any $v, v' \in \mathcal{V}$, such that $v(Y) = v'(Y)$, and for any $Y \in \mathcal{Y}$, we may insist that $\gamma(Y, v') = \gamma(Y, v)$. While all the results below hold for either restriction, we present the most general case in line with the rest of the literature.

3. *Overall-balanced* – $\sum_{i \in N} \gamma_i(Y, v) = v(Y)$ for all $v \in \mathcal{V}$ and $Y \in \mathcal{Y}$. Hence, the allocation rule should assign the full value of the network to the agents (Jackson and van den Nouweland, 2005).
4. *Component-decomposable* – $\gamma_i(Y, v) = \gamma_i(Y(S), v)$ for each $v \in \mathcal{V}$, $Y \in \mathcal{Y}$, $S \in \Pi(Y)$, and $i \in S$. The allocation of value in each component does not depend on the structure of other components (Jackson and van den Nouweland, 2005).

Taken together, assumptions 3 and 4 require that the value of any component is allocated only to the members of the component (Jackson and van den Nouweland, 2005).

2.3.2 Stability

We now extend the definition of strong stability (Dutta and Mutuswami, 1997) to networks represented by hypergraphs. We say that a group of agents $S \subseteq N$ *induces* a network $Y' \in \mathcal{Y}$ from $Y \in \mathcal{Y}$ if:

- New agreements are formed only by agents in S : $y \in Y'$ and $y \notin Y \implies y \subseteq S$
- Agents in S can sever any existing contractual agreement with those within or outside S : $y \in Y$ and $y \notin Y' \implies y \cap S \neq \emptyset$.

A group $S \subseteq N$ *deviates* from network Y with respect to an allocation rule γ and a value function $v \in \mathcal{V}$ if for all $i \in S$, $\gamma_i(Y', v) > \gamma_i(Y, v)$ where Y' is a network induced from Y by S .

Definition 2.1. A network $Y \in \mathcal{Y}$ is *strongly stable* with respect to an allocation rule γ and a value function $v \in \mathcal{V}$ if no group $S \subseteq N$ deviates from Y .

A strongly stable network is robust to deviations that leave all the deviators strictly better off. Jackson and van den Nouweland (2005) use an even stronger definition of

“strong stability” by ruling out any group deviations, in which all deviating agents are weakly better off and at least one agent is strictly better off. ¹¹

2.3.3 Key assumption about the allocation rule

We now formally state our final and crucial assumption on the allocation rule, which extends the definition of pairwise alignment in Pycia (2012) to hypergraphs.

Definition 2.2. For a given value function $v \in \mathcal{V}$, an allocation rule γ satisfies *strong pairwise alignment* if for any $Y, Y' \in \mathcal{Y}$, $C \in C(Y)$ and $C' \in C(Y')$ and any agents i and j such that $i, j \in C \cap C'$, we have that $\gamma_i(Y, v) \geq \gamma_i(Y', v) \iff \gamma_j(Y, v) \geq \gamma_j(Y', v)$.

The assumption states the following. For a given value function, pick any two networks Y and Y' . Then, pick any two components on these networks $C \in C(Y)$ and $C' \in C(Y')$. Suppose that there are some agents $S \subseteq N$ who are members of both components and there exists $i \in S$ whose payoff is higher in $C \in C(Y)$. Strong pairwise alignment states that the allocation rule should be such that all agents in S have higher payoff in $C \in C(Y)$. Note that we can restrict our attention to components of networks because the payoff to each agent depends only on the value of the component to which he belongs.

2.3.4 Stability result

Our first result states that strong pairwise alignment of the allocation rule is a necessary and sufficient condition for strong stability in networks.

Theorem 2.1. *A network $Y \in \mathcal{Y}$ is strongly stable with respect to the regular, anonymous, overall-balanced, and component-decomposable allocation rule γ and the value function $v \in \mathcal{V}$ if and only if the allocation rule satisfies strong pairwise alignment.*

¹¹There is a distinction between strong stability and the classic definition of the core (aside from the overlapping coalition structure). Strong stability allows agents to keep some of their existing links after a deviation; the core does not. Core allows arbitrary utility transfers within a coalition, whereas, in our definition of strong stability, payoffs are fully determined by the allocation rule. Shapley (1967) establishes necessary and sufficient condition for the existence of the core and Shapley (1971) shows that the core is always non-empty in cooperative games with a supermodular valuation function.

Theorem 2.1 echoes the key result by Jackson and van den Nouweland (2005), who characterised the existence of strongly stable networks (in their sense, see section 2.3.2) with an egalitarian allocation rule, which precludes any relationship between network position and share of component value. While the egalitarian allocation rule satisfies strong pairwise alignment, we permit a richer class of allocation rules that may, in fact, take the position and the role of an agent in the network into account. Example 1 highlights the limited sense in which payoffs can depend solely on an agent’s position in the network.

Example 1 Consider a very simple network in which $N = \{1, 2, 3\}$ and agents can only make bilateral links (i.e. $Y \subseteq \{y \in 2^N \mid |y| = 2\}$). The value of any empty component is 0; the value of any two-agent component is 2; the value of a line network is β ; value of the triangle is 6. Suppose that value is allocated in proportion to agents’ influence, measured by eigenvector centrality, on the component (Newman, 2010).¹² Then value is split equally in the triangle (2 each) and in the two-agent components (1 each). In the line component, the ends receive $\beta(1 - \frac{\sqrt{2}}{2})$ and the center receives $\beta(\sqrt{2} - 1)$. If $\beta = 4$, strong pairwise alignment holds and a strongly stable (and in this case efficient) triangle network exists. Interestingly, when $\beta = 5$, the triangle is still a strongly stable network, even though pairwise alignment no longer holds (the center node prefers the line to the triangle, but the ends do not, so they maintain their link). However, in this case, Theorem 2.1 does not apply as the allocation rule is not regular.

■

If we rule out within-component externalities, we can effectively retrieve the coalition formation framework of Pycia (2012). In order to illustrate this, let us simply assume that agents’ payoffs equal the sum of the values of the contracts they sign. We say

¹²This is a common measure of node influence in a network. Eigenvector centrality vector is the eigenvector corresponding to the largest eigenvalue λ that solves $\mathbf{A}\mathbf{x} = \lambda\mathbf{x}$, where \mathbf{A} is the adjacency matrix of the simple graph representing the network. These definitions can be extended to hypergraphs.

that the allocation rule is *y-additive* if $\gamma_i(Y, v) = \sum_{y \in Y_i} \gamma_i(y, v)$. This is equivalent to additive separability across contracts (Crawford and Knoer, 1981).

Corollary 2.1. *Suppose the allocation rule γ is y-additive. A network $Y \in \mathcal{Y}$ is strongly stable with respect to the regular, anonymous, overall-balanced, and component-decomposable allocation rule γ and the value function $v \in \mathcal{V}$ if and only if strong pairwise alignment is satisfied for any $y \in Y$ and $y' \in Y'$.*

As Pycia (2012) showed, equal sharing or Nash bargaining (with exogenous bargaining powers and fixed outside options) for value within each contract satisfies pairwise alignment. Assuming a *y-additive* allocation rule, we can immediately extend this claim by allowing agents to have different Nash bargaining powers across different contracts in a network. However, Example 2 below shows that bargaining according to Kalai and Smorodinsky (1975) rules in some contracts may lead to the overall failure of strong pairwise alignment and therefore to non-existence of a strongly stable network.

Example 2 This example extends examples presented in Pycia (2012) to a hypergraph and provides further intuition about the conditions under which strong pairwise alignment holds. In this example, we show that in a network without strong pairwise alignment of the allocation rule there will not necessarily be a strongly stable allocation. Let us suppose there are three venture capital funds $\{f_1, f_2, f_3\}$ and two inventors $\{g_1, g_2\}$, so $N = \{f_1, f_2, f_3, g_1, g_2\}$. Inventors can sign up to two contracts with funds because funds may specialise in different managerial fields. Production technology of funds is complementary in inventors: unless a fund can find at least two inventors, its output is zero because it cannot bring the invention to the market. We assume that the agents' utility is *y-additive*, so there are no within-component externalities. f_1 offers equal sharing of value, whereas f_2 and f_3 offer bargaining according to Kalai-

Smorodinsky shares. Utility functions are:

$$u_{f_1}(s) = u_{g_2}(s) = u_{f_3}(s) = s^{1/2}$$

$$u_{f_2}(s) = s$$

$$u(s) = s^{1/6}$$

Contract values are:

$$v(\{f_1, g_1, g_2\}) = v(\{f_2, g_1, g_2\}) = 99$$

$$v(\{f_3, g_1, g_2\}) = 42$$

$$v(\{\cdot, g_1\}) = 44$$

$$v(\{\cdot, g_2\}) = 64$$

Otherwise $v(\cdot) = 0$. The resulting payoffs for each agent in each contract are reported in Table 2.1.

Table 2.1: Payoffs in a network that is not strongly stable

Fund	Inventors		
	$\{g_1, g_2\}$	$\{g_1\}$	$\{g_2\}$
f_1 (equal sharing)	33,33,33	22,22	32,32
f_2 (Kalai-Smorodinsky)	59,5,35	30,14	32,32
f_3 (Kalai-Smorodinsky)	19,4,19	30,14	32,32

Comment: Output shares are listed in the following order: fund, respective inventors

If inventors g_1 and g_2 form a contract with fund f_1 , then they might also have an incentive to join fund f_2 . But in that case, f_3 and g_1 would have an incentive to deviate.

Hence,

$$\begin{aligned} \gamma_{f_3, g_1}(\{\{f_1, g_1, g_2\}, \{f_2, g_1, g_2\}, \{f_3\}\}) &< \gamma_{f_3, g_1}(\{\{f_1, g_1, g_2\}, \{f_3, g_1\}, \{f_2, g_2\}\}) \\ \gamma_{f_3, g_2}(\{\{f_1, g_1, g_2\}, \{f_3, g_1\}, \{f_2, g_2\}\}) &< \gamma_{f_3, g_2}(\{\{f_1, g_1, g_2\}, \{f_3, g_2\}, \{f_2\}\}) \\ \gamma_{f_2, g_2}(\{\{f_1, g_1, g_2\}, \{f_3, g_2\}, \{f_2\}\}) &< \gamma_{f_2, g_2}(\{\{f_1, g_1, g_2\}, \{f_2, g_1, g_2\}, \{f_3\}\}) \end{aligned}$$

and hence there is no strongly stable network. ■

Example 2 makes it clear how the results from Pycia (2012) translate into a network structure. If the allocation rule is y -additive, failure to satisfy pairwise alignment across individual contracts translates into failure to satisfy strong pairwise alignment. In Example 2 above, it is worth noting that strong pairwise alignment fails not only because of Kalai-Smorodinsky bargaining does not satisfy pairwise alignment in components containing agents f_2 or f_3 , but also because these two agents offer a different bargaining procedure from f_1 . Hence, because we have that $\gamma_{g_1}\{f_2, g_1, g_2\} < \gamma_{g_1}\{f_1, g_1, g_2\}$, but $\gamma_{g_2}\{f_2, g_1, g_2\} > \gamma_{g_2}\{f_1, g_1, g_2\}$, strong pairwise alignment also fails.

2.3.5 Efficiency

Dutta and Mutuswami (1997) carefully examined the conflict between stability, efficiency, and symmetry in strongly stable networks. In particular, they found that there exists no allocation rule with certain desirable properties which guarantees that any efficient network is strongly stable. The example below shows that under strong pairwise alignment, strongly stable networks may not be efficient.

Example 3 We now add another inventor g_3 to agents described in Example 2, so now $N = \{f_1, f_2, f_3, g_1, g_2, g_3\}$. The roles the agents as venture capital funds and inventors remain as above. There is Nash bargaining (in which the outside options give zero utility) over profits in all contracts, and agents can sign at most two contracts.

Agents' utility functions are:

$$u_{f_1}(s) = 2s^{1/2}$$

$$u_{f_2}(s) = s^{1/3}$$

$$u_{f_3}(s) = s^{1/4}$$

$$u_{g_l}(s) = s^{1/l+4}$$

for $l = 1, 2, 3$, where s is the share of profits from a contract. Contract values are:

$$v(\{f_1, g_1, g_2, g_3\}) = 80 \quad v(\{f_1, g_i, g_j\}) = 50$$

$$v(\{f_2, g_1, g_2, g_3\}) = 90 \quad v(\{f_2, g_i, g_j\}) = 40$$

$$v(\{f_3, g_1, g_2, g_3\}) = 70 \quad v(\{f_3, g_i, g_j\}) = 60$$

Otherwise $v(\cdot) = 0$. The resulting payoffs for each agent in each contract are reported in Table 2.2.

Table 2.2: Inefficiency in a strongly stable network

Fund	Inventors			
	$\{g_1, g_2\}$	$\{g_2, g_3\}$	$\{g_1, g_3\}$	$\{g_1, g_2, g_3\}$
f_1	28.9, 11.5, 9.6	30.9, 10.3, 8.8	29.7, 11.9, 8.5	24.8, 9.9, 8.3, 7.9
f_2	19.1, 11.4, 9.5	20.7, 10.4, 8.9	19.7, 11.8, 8.5	35.6, 21.4, 17.8, 15.3
f_3	24.3, 19.5, 16.2	26.8, 17.9, 15.3	25.3, 20.2, 14.5	19.8 15.8 13.2 11.3

Comment: Output shares are listed in the following order: fund, respective inventors. **Bold:** strongly stable network

The efficient network is $\{\{f_1, g_1, g_2, g_3\}, \{f_2, g_1, g_2, g_3\}\}$. In this example, the unique strongly stable network $\{\{f_2, g_1, g_2, g_3\}, \{f_3, g_2, g_3\}\}$ is not efficient. No agent wants to contract with fund f_1 because it has too much bargaining power, and inventors g_2 and g_3 “use” inventor 1 to obtain the highest value. ■

2.4 Multilateral matching markets

We now extend the network model to a matching market in which there are several types of agents. The examples in the previous section presented a two-sided model with venture capital funds and inventors. We now allow the market to have more than two types of agents, but agents can match with at most one other agent on their own side of the market in any given contract. For example, in order to produce profitable ventures, inventors must be matched to funds and universities. Each fund can enter into several ventures with inventors and universities, but each venture will require a separate contractual agreement. We allow agents to have general preferences over their contractual agreements and partners. In the network model, for a given value function, the continuous allocation rule could specify infinitely many profit levels for any agent. In the multilateral matching markets model, we assume that contract terms come from a finite set.

2.4.1 Ingredients

We now introduce a multilateral matching model with contracts between agents in N .¹³ Let X denote the set of all contracts between agents in N . A *contract allocation* is a non-empty set of contracts $Y \subseteq X$. Each multilateral contract $y \in Y \subseteq X$ is associated with several types of agents (such as funds and inventors) and includes contract terms. Denote by $a(y) \subseteq N$ the set of agents in contract y . If $f \in a(y)$ then the fund f is associated with a contract y . Then $a(Y) \equiv \cup_{y \in Y} \{a(y)\}$ denotes the set of agents who have contracts in Y . In order to preserve a multipartite structure, we make the following assumptions on contractual *feasibility*:

1. Agents can be partitioned into at least two sides of the market: $N = \{N_m, \dots, N_M\}$

¹³Hatfield and Kominers (2011b) also consider a multilateral matching model, but they focus on the existence of competitive equilibrium in a transferable utility setting with quasilinear utility functions.

where $N_m \cap N_n = \emptyset \forall m, n = 1, 2, \dots, M$ such that $m \neq n$ and $M \geq 2$.

2. There are at least two agents on each side: $|N_m| \geq 2, \forall m = 1, 2, \dots, M$.
3. Any single contract contains at most two agents on the same side: $|a(y) \cap N_m| \leq 2$ for all $y \in X$ and $\forall m = 1, 2, \dots, M$.

For a three-sided market, we could partition the finite set of agents $N = F \cup G \cup H$ into a set F of venture capital funds, a set G of inventors, and set H of universities, but generally there can be an arbitrary number of partition elements (see Figure 2.3). For example, we could set $X \equiv F^2 \times G^2 \times H^2 \times \Theta$, where a finite Θ specifies contract terms.¹⁴ Note that in contrast to the network model, we now allow the same set of agents to sign several contracts. Additionally, denote $Y_i \equiv \{y \in Y | i \in a(y)\}$ as the set of contracts in Y associated with agent i . Therefore, $a(Y_i)$ are the agents with whom agent i has contracts in Y (including i).¹⁵ We will denote by $Y_{i,j}$ any Y such that $\exists y \in Y$ for which $\{i, j\} \subseteq a(y)$ i.e. any contract allocation in which i and j are members of at least one contract.

Define \succsim_i as weak preference relation (a complete preorder) over sets of contracts in 2^{Y_i} involving i . For a set of contracts Y_i in the set of contracts X , this preference relation is denoted $\succsim_i \in \mathcal{R}_i^X$ where \mathcal{R}_i^X is the domain of all possible preference orderings on X for i . The preference profile of all agents over X is then $(\succsim_i)_{i \in N} \equiv \succsim_N$. Hence, \succsim_N is an element of $\mathcal{R}^X \equiv \prod_{i \in N} \mathcal{R}_i^X$ – the preference domain for a set of contracts X .

We can now reintroduce strong pairwise alignment in the context of multilateral matching markets.

Definition 2.3. A preference profile $\succsim_N \in \mathcal{R}^X$ satisfies *strong pairwise alignment* if for any agents $i, j \in N$ and any contract allocations $Y, Y' \subseteq X$ such that $\{i, j\} \subseteq$

¹⁴Note, that we allow agents to sign contracts which contain at most one other agent on the same side of the market.

¹⁵ $|a(Y_i)| = 1$ represents an unmatched agent.

$a(Y) \cap a(Y')$ and $Y_i \cap Y_j \neq Y'_i \cap Y'_j$, we have that $Y \succsim_i Y' \iff Y \succsim_j Y'$.

There are some subtle differences between this definition and the definition of strong pairwise alignment over networks. In the multilateral matching market model, agents' preferences are defined only over contracts to which they belong. While there may be externalities within and across their contracts, agents are not affected by the contracts signed by other agents. However, in the networks model, we permitted agents' payoffs to be affected by the contracts of others since the allocation rule determined payoffs of all agents in a given component. Yet the basic intuition remains the same. Whenever there is a choice between two different sets of contracts Y and Y' , both of which contain different contracts that include i and j , i prefers Y to Y' if and only if j does. Loosely speaking, if funds and inventors are partners in the project, they aim to maximise the value of the project. Similarly, if an inventor wants to sign a contract for a venture that pursues a certain management strategy, then the fund in the venture would support her. The following example shows how we operationalise the definition.

Example 4 Suppose there are seven agents $N = \{i_1, \dots, i_7\}$, and we want to check for strong pairwise alignment between two contract allocations. Suppose one contract allocation is $Y = \{\omega, \phi, \chi, \psi\}$, where $a(\omega) = \{i_1, i_2\}$, $a(\phi) = \{i_2, i_3\}$, $a(\chi) = \{i_1, i_3\}$ and $a(\psi) = \{i_1, i_4, i_5\}$, and another is $Y' = \{\eta, \psi, \kappa\}$, where $a(\eta) = \{i_1, i_2, i_3\}$ and $a(\kappa) = \{i_6, i_7\}$. Strong pairwise alignment would require that

$$Y \succsim_{i_1} Y' \iff Y \succsim_{i_2} Y' \iff Y \succsim_{i_3} Y'$$

We do not impose any restrictions on the preferences between i_1, i_4 and i_5 (because they were party to the same contract in both contract allocations) or on the preferences of $\{i_6, i_7\}$ (because they were not members of Y). ■

2.4.2 Heterogeneous contracts

In the multilateral matching market, we no longer assume that contracts simply specify shares of perfectly divisible profit. Instead, we assume that agents specify contract terms for many dimensions of their relationship. Contract terms are necessarily finite. The rationale for this was given by Roth (1984b, p. 49):

elements of a job description can take on only discrete values; salary cannot be specified more precisely than to the nearest penny, hours to the nearest second, etc.

In our context, a contract could stipulate exactly how much equity each party gets, whether it gets a seat on the management board, and so on. In this chapter, we will assume that contractual terms are sufficiently different across contracts.¹⁶ Following Pycia (2012), we impose two conditions on the preference domain over the set of contracts for all agents. Then we interpret these conditions in terms of heterogeneity of the contract set.

Definition 2.4. A contract set X is *heterogeneous* if, for any $\succsim_N \in \mathcal{R}^X$, for any agent $i \in N$, and for any three different sets of contracts $Y_i, Y'_i, Y''_i \subseteq X$ for i , if $Y'_i \succsim_i Y''_i$ then there is a preference profile $\succsim^*_N \in \mathcal{R}^X$ such that $Y'_i \succ^*_i Y_i \succ^*_i Y''_i$, and all agents' \succsim^*_N -preferences between sets of contracts not including Y are the same as their \succsim_N -preferences.

Definition 2.5. A contract set X is *very heterogeneous* if it is heterogeneous and:

- For any $\succsim_N \in \mathcal{R}^X$, and for any two different sets of contracts $Y_{i,j}, Y'_{i,j} \subseteq X$ for $\{i, j\}$, there is a preference profile $\succsim^*_N \in \mathcal{R}^X$ such that $Y_{i,j} \prec^*_{i,j} Y'_{i,j}$, and all agents' \succsim^*_N -preferences between sets of contracts not including Y are the same as their \succsim_N -preferences.

¹⁶We can embed a theory of contractual language developed by Hatfield and Kominers (2011a) (which extends the work of Roth, 1984b) in our model of a multilateral matching market. We discuss the relationship between our results in section 2.4.4.

- For any $\succsim_N \in \mathcal{R}^X$, for any agents i and j and for three different sets of contracts $Y_{i,j}, Y'_i, Y''_j \subseteq X$ for i and j , if $Y'_i \prec_i Y_{i,j} \sim_j Y''_j$, then there is a preference profile $\succsim^*_N \in \mathcal{R}^X$ such that $Y'_i \prec^*_i Y_{i,j} \prec^*_j Y''_j$, and all agents' \succsim^*_N -preferences between sets of contracts not including Y are the same as their \succsim_N -preferences.

Contracts could be heterogeneous for a number of reasons. First, contracts could specify many dimensions of a relationship. Therefore, a consistent ranking of contract sets (even of those involving the same agents) across the entire preference domain is unlikely. This contrasts the approach of Kelso and Crawford (1982), who assume that workers rank all contracts with a particular firm according to the wage. Second, contracts could be contingent on some unknown state of the world that could be affected by other agents in the contract. The following example illustrates this point. Suppose an inventor i wants to sign a contract for a venture with one of three funds. Contract Y' gives the inventor less equity than contract Y'' . We want to justify i 's preference profile $Y' \succsim_i Y \succsim_i Y''$ in the case where Y gives him *even less* equity than Y' . The inventor's payoff is contingent on the success of his product. Hence, inventor's payoffs could depend on the level of equity and (exogenous) quality of the fund manager. In fact, inventors regularly turn down "better" contracts for the opportunity to work with more experienced fund managers.

More generally, a contract set is heterogeneous if there exists a preference profile such that any agent who is party to a contract set Y is indifferent between it and any Y' and Y'' , while leaving the preferences of others unchanged. A contract set is very heterogeneous if agents on an intersection of two contract sets all strictly prefer one contract set to another in some preference profile. While most of the matching literature assumes that all preference relations between contracts are strict, our conditions on preferences also allow for indifferences.¹⁷ Contracts that simply specify division of profit

¹⁷Erdil and Kumano (2011) also consider a matching market with weak preferences.

according to an appropriately defined continuous and strictly increasing allocation rule – as in the networks model – are very heterogeneous (because they do not need to specify the value function). A universal preference domain over contracts means that the contract set is very heterogeneous. This extreme case may, for example, represent informal or oral contract agreements.

2.4.3 Stability

We say that a contract allocation is strongly stable if no agent wants to drop his contracts (individual rationality); and if no group of agents can deviate, form any new contracts only among themselves, drop any of their existing contracts and strictly prefer the new allocation. The definition of strongly stable contract allocations mimics the strong stability notion we used for networks. Since most of the matching literature deals only with strict preferences, it is the strongest stability concept we can define. If a strongly stable contract allocation exists, then it is difficult to imagine how any deviation from it could be rational.¹⁸

¹⁸Hatfield and Kominers (2011a) call our strong stability concept *strong group stability*, but we keep our nomenclature consistent with the networks model presented in this chapter. Their definition of “strong stability” in two-sided markets does not require that the contract allocation be individually rational. Weaker stability concepts that have been proposed in the many-to-many (two-sided) matching market literature include:

1. *Stability* (Hatfield and Kominers, 2011a) is strong stability with the additional requirement that the new set of contracts must be optimal, not merely an improvement.
2. *Group stability* (Roth and Sotomayor, 1990, Konishi and Ünver, 2006) is strong stability with the additional requirement that the blocking agents must agree which contracts to keep after the deviation.
3. *Setwise stability* (Sotomayor, 1999) is strong stability with the requirements that (i) blocking agents must agree which contracts to keep after deviation *and* (ii) that the blocking set of contracts is individually rational. In the context of a two-sided market, Echenique and Oviedo (2006) showed that substitutability on one side of the market and *strong* substitutability (if an option is accepted, then it is accepted from a *worse* choice set) on the other side of the market is sufficient for setwise stability. Klaus and Walzl (2009) extended these results to a contractual setting.
4. *Core* (*corewise stability*) is strong stability with the additional requirement that blocking agents must drop all existing contracts and may only form new contracts among themselves.
5. *Many-to-one stability* (Hatfield and Milgrom, 2005, Hatfield and Kominers, 2011a) is setwise stability with the additional requirement that there is only one agent on one side of the market in the blocking set.

Definition 2.6. A contract allocation $Y \subseteq X$ is *strongly stable* if:

1. It is *individually rational*: for all $i \in a(Y)$, $Y_i = \max_{\succsim_i} \{U \subseteq Y \mid y \in U \Rightarrow i \in a(y)\}$.
2. There does not exist a non-empty, feasible *blocking set* of contracts $Z \subseteq X$ such that:
 - $Z \cap Y = \emptyset$, and
 - For all $j \in a(Z)$, there exists a $W_j \subseteq Z \cup Y$ such that $Z \subseteq W_j$ and $W_j \succ_j Y_j$.

We call a contract allocation that is induced by the failure of individual rationality or by a blocking set of contracts a *profitable deviation*. As we rule out any reasonable deviation in a strongly stable contract allocation, we can construct a preference profile such that any contract allocation can be induced from any other contract allocation with just one profitable deviation. This turns out to be a key step to extend Pycia's results to a multilateral matching market.

2.4.4 Main results

The two other main results of this chapter are the following:

Theorem 2.2. *If contract set X is heterogeneous and all preference profiles in \mathcal{R}^X satisfy strong pairwise alignment, then a strongly stable contract allocation $Y \subseteq X$ exists in all preference profiles in \mathcal{R}^X .*

Theorem 2.3. *If contract set X is very heterogeneous and a strongly stable contract allocation $Y \subseteq X$ exists in all preference profiles in \mathcal{R}^X , then all preference profiles in \mathcal{R}^X satisfy strong pairwise alignment.*

6. *Pairwise stability* is setwise stability with the additional requirement that the size of the blocking set is 2. Roth (1984b) showed that substitutability is sufficient for a pairwise stable allocation in a many-to-many matching market with contracts.

To prove the first result, we first show that if there are no profitable deviation cycles, then a strongly stable contract allocation exists. Then we show that if strong pairwise alignment holds, then there are no profitable deviation cycles of size three. Finally, we show that absence of these cycles implies absence of cycles of any size. To prove the second result, we show that strong stability implies absence of profitable deviation cycles of size three. Absence of these cycles turns out to guarantee strong pairwise alignment in every preference profile.

These results show that a strongly stable contract allocation can exist in a multilateral matching market with finite contracts despite complementarities and externalities as long as the agents' preference profiles are constrained by strong pairwise alignment. There is interesting tension between the strong pairwise alignment assumption and the assumptions on the heterogeneity of contracts. In every preference profile, agents' preferences over sets of contracts need to be sufficiently aligned *within* the preference profile. Therefore, there cannot exist any two contracts, which drastically differ in the relative favourability of contract terms. However, as we explained above, contracts need to be sufficiently heterogeneous in order to accommodate the different rankings *across* preference profiles. This heterogeneity is achieved precisely by the existence of multidimensional contract terms.

Our setting puts several recent results in the matching literature in perspective. In a many-to-many two-sided matching market with contracts (e.g. between doctors and hospitals), Hatfield and Kominers (2011a) show that if preferences of one agent are not substitutable in some preference profile, then “there exist substitutable preferences for the other doctors and hospitals such that no many-to-one stable allocation exists” in another preference profile (Theorem 12).¹⁹ Our results do not contradict this theorem for two reasons. First, we assumed that the matching market is multilateral

¹⁹This result also implies that no strongly stable allocation exists.

and imposed different conditions on contractual feasibility. Second, in Theorem 2.2, we assumed that all preferences profiles in the domain satisfy strong pairwise alignment, while Hatfield and Kominers (2011a) do not. However, we showed that substitutability of preferences is certainly not necessary (in the usual mathematical sense) for strong stability in a multilateral matching market.

Hatfield and Kominers (2011a) also show that a “coarser contractual language” will support a strongly stable many-to-many matching whenever the coarsening does not remove the strongly stable allocation (Theorem 3). In their framework, a coarser contractual language effectively precludes the existence of some contracts by removing the means to express them. For a given substitutable preference profile, removing contracts preserves substitutability over the remaining contracts in the new profile and continues to guarantee strongly stable contract allocations. Therefore, there ought to be a trade-off between coarseness of the contractual language and stability. Our condition on the heterogeneity of the contract set guarantees a large domain of preference profiles rather than simply increasing the size of the contract set. If we coarsen the contractual language in the sense of Hatfield and Kominers (2011a), we are unable to express certain contractual terms. Hence, contracts contain fewer dimensions and our assumptions on contract heterogeneity are less likely to be satisfied. While the strongly stable contract allocation may still be available to the agents, there may not be a preference profile to support it. Thus, the spirit of our results suggests that heterogeneous contracts and strong pairwise alignment are complementary forces in guaranteeing strongly stable contract allocations in multilateral matching markets. If contractual language is expressive in the sense of Hatfield and Kominers, then preferences over contracts in our model written in that contractual language are more likely to be rich. What we show is that without sufficient variation in contract terms, strongly stable contract allocations are not guaranteed to exist even when strong pairwise alignment is satisfied.

2.5 Applications of the models

The following three examples illustrate applications of the two models presented in this chapter.

2.5.1 Production networks

In *The Wealth of Nations*, Adam Smith estimated that in a pin factory ten specialised workers produced 4,800 pins a day each, whereas any one of them working alone would produce

certainly, not the two hundred and fortieth, perhaps not the four thousand eight hundredth part of what they are at present capable of performing, in consequence of a proper division and combination of their different operations (Smith (1776), Book 1, Chapter 1)

Modern production of sophisticated technology involves the assembly of various complementary components, which may be manufactured by different firms around the world. Ostrovsky (2008, p. 914) notes that:

In some industries (e.g., construction), firms along supply chains combine several complementary inputs to produce final goods, with inputs themselves consisting of multiple complementary parts, many of them heterogeneous, complex, and an important part of the final cost of the outputs.

Using a two-sided matching model, Fox (2009) estimated production functions of automotive suppliers and assemblers. He found that there are substantial complementarities for parts suppliers in matching with one assembler headquartered in Asia and another outside Asia. Similarly, venture capital funds provide complementary inputs of capital and managerial experience across a portfolio of start-ups, which are encouraged to specialise and share each other's technologies (Hellman, 1998). Contractual arrangements in ventures play an important role in investment decisions (Macho-Stadler et al., 2008) because when firms innovate and specialise, their technologies “spill over” as their

competitors imitate and improve them. Finally, firms may make their R&D decisions depending on their position in the production network (Baker et al., 2008).

2.5.2 Teams

Peer effects are externalities that arise due to social interaction. Workers in a firm may be assigned to work simultaneously in several teams, and their productivity may be affected by the composition of their team and the structure of the organisation (Demange, 2004). Page Jr. and Wooders (2010) introduced an interesting application of this in a non-cooperative setting. In their model, agents join multiple clubs (that provide local public goods) and multiple activities within these clubs. Baccara et al. (2012) provided a fascinating empirical study of how overlapping networks – social and work-related – may account for the efficiency and stability of a matching market (of faculty members into offices). Our model allows for the presence of peer effects in matching markets and externalities in networks and therefore can be applied to study these settings.

2.5.3 Firm-union bargaining

A common criticism of demanding stability concepts is that group deviations are hard to achieve since they require substantial coordination on behalf of the deviating agents. Most market design applications focus on pairwise stable allocations – those immune to deviations by pairs of agents (Roth, 1984b). However, a natural setting where group deviations may occur is in a bargaining process between firms and unionised workers (or workers, who are covered by a collective bargaining agreement and are not members of a union). The union acts as a credible coordinator of arbitrary deviations by workers: for example, it can ballot workers to go on strike. In this context, the set of pairwise stable outcomes may be an imprecise predictor of market outcomes whereas strongly stable outcomes are likely to be more accurate. The multilateral market model

presented in this chapter allows firms to hire both unionised and non-unionised workers and allows workers to take on jobs outside the remit of their union. We characterised the types of contracts that permit contract allocations that are robust to arbitrary deviations of unionised workers. Since Nash bargaining solution turned out to be central to stability and our results provide a new justification for its use in models of firm-union bargaining (Oswald, 1982, Grout, 1984).

2.6 Conclusion and extensions

In this chapter, we showed that strong pairwise alignment is a necessary and sufficient condition for the existence of strongly stable networks and contract allocations in multilateral matching markets. First, this condition generalises some results from the networks literature. Second, we showed that substitutability of preferences is not essential for the existence of strongly stable contract allocations in multilateral matching markets. We also offered a new interpretation of the relationship between contract heterogeneity and production technologies in matching markets. This chapter shows that an integrated approach to matching markets and networks can be fruitful in generating new insights.

An important practical problem is to design an efficient algorithm that would test whether strong pairwise alignment is satisfied. Several papers have proposed algorithms to find stable allocations whenever they exist in many-to-one matching markets (Echenique and Oviedo, 2004, Echenique and Yenmez, 2007), but it is not obvious whether they can be applied to the multilateral market model presented in this chapter. If an (efficient) algorithm does not exist, then one ought to consider whether the stable outcomes can be reached via random paths (Kojima and Ünver, 2008). Imposing conditions for stability on the preference profile makes strategy-proofness of a possible contract allocation mechanism difficult to interpret, however, it is easy to show that any

strongly stable outcome in the network or the matching market model can be supported by a Strong Nash equilibrium in an appropriately defined strategic game (see, Dutta and Mutuswami (1997) and Pycia (2012)). Finally, in contrast to Hatfield and Kominers (2011b, Theorem 6), Example 2 showed that strongly stable outcomes in matching markets or networks are not necessarily efficient under strong pairwise alignment. This conflict between strong stability and efficiency, highlighted by Dutta and Mutuswami (1997), could be explored further.

2.7 Appendix

Proof of Theorem 2.1. First, note that because strong pairwise alignment holds for any two components $C \in C(Y)$ and $C' \in C(Y')$ that contain the same set of agents, the preferences of these agents over the components must be same. Therefore, no network is strongly stable if it contains a component which has a strictly lower value than a component containing the same agents in another network. For any set of agents $S \subseteq N$ and a given v , let us define the “top” component(s) $K(S) = \arg \max_{C \in C(S)} \{v(C) | \Pi(S) = \{S\}\}$. Non-uniqueness of top components implies that if there are several components with identical value, then agents would be indifferent between them so in general strongly stable networks will not be unique. The structure of the problem is now isomorphic to the coalition formation setting in Corollary 1 in Pycia (2012). We can now apply Corollary 1 in Pycia (2012) over $K(S)$, which completes the proof. \square

2.7.1 Sufficiency of strong pairwise alignment for a strongly stable contract allocation

Remark The proofs here closely follow the logic and method of proofs of Theorems 1 and 2 in Pycia (2012), while making the necessary adjustments for the structure of contract allocations and stability definitions. There are two observations, which allow us to extend Pycia’s proof.

1. Any contract allocation can be reached from any other contract allocation within one profitable deviation (see definition in the section 2.4.3) by selecting an appropriate preference profile.
2. Assumptions about contract feasibility allow for the existence of a contract between any pair of agents, which is sufficient for Lemmata 2.3, 2.4, and 2.6.

Notation Since the set of agents $\{i, \dots, j\} = N$ is finite and $|N| = n$, we can relabel all agents so $N = \{i_1, i_2, \dots, i_n\}$. Without loss of generality, we can refer to agent 1, 2, 3... etc. Thus for any contract allocation Y , Y_1 denotes the contracts in Y to which i_1 is party. We denote by $Y_{1,2}$ any contract allocation such that for some $y \in Y_{1,2}$ $a(y) \ni i_1, i_2$. Note that $Y_{1,2}$ does not need to identify a unique contract allocation, rather it is a property of a particular contract allocation.

Define $\mathcal{Y} = \{Y \subseteq X \mid \forall y \in Y, a(y) = N\}$ as the set of contract allocations, such that in every allocation every contract contains all agents. This is a special corner case in which there are exactly two agents on each side of the market.

Lemma 2.1. *Let \succsim_N be a preference profile, such that the contract allocation $Y \subseteq X \setminus \mathcal{Y}$ is strongly stable and let \succsim'_N be a preference profile, which is identical across all sets of contracts $Y_i \succsim_i Y'_i \iff Y_i \succsim'_i Y'_i$ for all $i \in a(Y), a(Y')$ such that $Y, Y' \subseteq X \setminus \mathcal{Y}$. Then, either Y or some $Z \in \mathcal{Y}$ is strongly stable under \succsim'_N .*

Proof. Take any contract allocation Y that is strongly stable with respect to a preference profile \succsim_N . Because \succsim'_N is identical on all Y except $Z \in \mathcal{Y}$, Y is individually rational and can only be blocked by the agents joining some of the contracts in Z under the preference profile \succsim'_N . Hence, Y is either strongly stable under the preference profile \succsim'_N or a non-empty blocking set of contracts $Z \subseteq \mathcal{Z}$ blocks Y . In the latter case, all agents must prefer Z to their current contract allocation Y_i

$$Z_i \succ_i Y_i \quad \forall i \in N \tag{2.1}$$

If no Z is strongly stable under \succsim'_N , it must be either because

1. no Z is individually rational and some agents would prefer to drop their contracts.

Since Z is finite, there must be at least one agent who prefers to drop all contracts.

2. and/or every Z is blocked by another contract. If every Z were blocked by another contract allocation in \mathcal{Z} , then there must exist a strongly stable contract allocation in \mathcal{Z} (because every block is done by all agents), which leads to a contradiction. Therefore, there must be a chain of blocking contract allocations in \mathcal{Z} that is eventually blocked by a contract allocation outside \mathcal{Z} : $Z_i^1 \prec_i \dots \prec_i Z_i^M \forall i \in N$ and $D_i \succ Z_i^M$ for some $i \in a(D)$.

In either case, we have that $D_i \succ'_i Z_i \succ'_i Y_i$ for some $i \in N$ (in Case 1, it possible that $|a(D_i)| = 1$ if i drop his contracts with others). But, since $D_i \not\subseteq \mathcal{Z}$, the preferences of \succsim_i coincide with \succ'_i , we have that $D_i \succ_i Y_i$ for $i \in a(D)$ contradicting the strong stability of Y . \square

Notation A k -cycle is a configuration of contract allocations $Y_{i_1}, \dots, Y_{i_k} \subseteq X$ and agents $\{i_1, \dots, i_k\}$, such that $Y_{i_j} \succ_{i_j} Y_{i_{j+1}}$ (modulo k) for $j = 1, \dots, k$ where at least one preference is strict and every preferred element of the configuration is a profitable deviation.

Remark Failure of strong stability is caused either by failure of individual rationality and/or by blocking contracts. In the former case, “blocking members” are necessarily members of the strongly unstable contract set in the allocation. In the latter case, (by definition of strong stability) some blocking members may come from outside the strongly unstable contract set in the allocation. However, if a contract set $Z = X \setminus Y$ is blocked only by members of $N \setminus a(Y)$, then $a(Y)$ are not affected by this.

Lemma 2.2. *If there are no k -cycles for any $k = 3, 4, \dots$, then there is a strongly stable contract allocation on X .*

Proof. Proceed by induction on $|N|$. The claim is true for $|N| = 1$ as the agent simply picks his most preferred contract. Suppose* the claim is true whenever there are fewer agents than $|N|$. This means that there is a strongly stable contract allocation on

$X \setminus \mathcal{Y}$ by Lemma 2.1 (otherwise we iterate profitable deviations \mathcal{Y} as in Case 2 of Lemma 2.1 giving us the result). Towards a contradiction, assume there is no strongly stable contract allocation on N (i.e. possibly involving any $Y \subseteq X$). Then any contract allocation $Y \subseteq X$ is either:

1. not individually rational
2. and/or Y is blocked
 - (a) by (some) agents in $a(Y)$
 - (b) by (some) agents outside $a(Y)$

and profitable deviation resulting from Y is $Y' \subseteq X \setminus \mathcal{Y}$. If there is no Y' that can result from such a deviation, then Y forms a strongly stable contract allocation for $a(Y)$ alongside a strongly stable contract allocation that includes only agents $N \setminus a(Y)$ (by supposition* above). Hence, there is a profitable deviation in X for any $Y \subseteq X$. Since X is finite, there must be a k -cycle of size at least three, since a profitable deviation Y' from Y cannot be blocked by Y itself. Hence, we have a contradiction. \square

Lemma 2.3. *If contract set X is heterogeneous and all preferences profiles in \mathcal{R}^X satisfy strong pairwise alignment, then no preference profile in \mathcal{R}^X admits a 3-cycle.*

Proof. Consider agents $i_1, i_2, i_3 \in N$.

We assume towards a contradiction that there exists a strongly pairwise aligned preference profile \succsim_N with contract allocations such that:

$$Y_{3,1} \prec_1 Y_{1,2} \succsim_2 Y_{2,3} \succsim_3 Y_{3,1} \tag{2.2}$$

Our assumptions on contract feasibility allow us to find a contract allocation Y' that contains at least one other agent $i_4 \in N$, who has a contract with i_1 and i_2 , with

i_2 and i_3 , and with i_1 and i_3 (possibly with all of them in one contract). If all three contract allocations $Y'_{1,2}, Y'_{2,3}, Y'_{3,1}$ are the same for i_1, i_2 and i_3 , then the analysis still holds. Suppose, all these sets of contracts are different. Then the heterogeneity of the contract set allows us to find a preference profile $\succsim'_N \in \mathcal{R}^X$ such that:

$$Y_{1,3} \succsim'_3 Y'_{1,2} \succsim'_1 Y_{1,2} \tag{2.3}$$

$$Y_{1,2} \succsim'_2 Y'_{2,3} \succsim'_2 Y_{2,3} \tag{2.4}$$

$$Y_{2,3} \succsim'_3 Y'_{1,3} \succsim'_3 Y_{1,3} \tag{2.5}$$

By strong pairwise alignment of \succsim'_N , we have from (2.5) that $Y'_{1,3} \succsim'_1 Y_{1,3}$, and from (2.3) that $Y_{1,3} \succsim'_1 Y'_{1,2}$, so by combining these preferences we obtain that:

$$Y'_{1,3} \succsim'_1 Y'_{1,2}$$

$$\text{from (2.3) and (2.4): } Y'_{1,2} \succsim'_2 Y'_{2,3}$$

$$\text{from (2.4) and (2.5): } Y'_{2,3} \succsim'_3 Y'_{1,3}$$

and at least one must be strict as we assumed that $Y_{1,3} \prec_1 Y_{1,2}$. By strong pairwise alignment $Y'_{1,3} \succsim_4 Y'_{1,2} \succsim_4 Y'_{2,3} \succsim_4 Y'_{1,3}$ with at least one preference relation being strict, which contradicts the initial assumption. \square

Lemma 2.4. *If contract set X is heterogeneous and if no preference profile in \mathcal{R}^X admits a 3-cycle, there no preference profile in \mathcal{R}^X admits an n -cycle.*

Proof. This lemma does not require strong pairwise alignment to hold.

We proceed by induction. For an inductive argument, we want to show that if no

preference profile admits a $3, \dots, m - 1$ cycle, then no preference profile in \mathcal{R}^X admits an m -cycle ($m > 3$). Towards a contradiction, let us assume some $\succsim_N \in \mathcal{R}^X$ admits an m -cycle

$$Y_{m,1} \succsim_1 Y_{1,2} \succsim_2 Y_{2,3} \dots \succsim_m Y_{m,1} \quad (2.6)$$

with at least one preference strict. Without loss of generality, we can assume that all i are different. By contractual feasibility, there exists some contract allocation Y in which i_1 and i_3 are parties to at least one contract (i.e. $i_1 \in a(Y_3)$).

There are two possible cases. In the first, Y is a part of the prohibited m -cycle. In the second, Y is not a part of the prohibited m -cycle.

1. $Y = Y_{l,l+1}$ for at least one $l = 1, \dots, m$. This means that Y also contains a contract to which i_l and i_{l+1} are party. Let us focus on $Y_{1,2}, Y_{2,3}, Y$. According to the antecedent of the lemma, $Y_{1,2}, Y_{2,3}$ and Y must have no 3-cycle. More concretely, we assumed that $Y_{1,2} \succsim_2 Y_{2,3}$, therefore, $Y_{1,2} \succsim_2 Y_{2,3} \succsim_3 Y \succsim_1 Y_{1,2}$ with at least one preference strict cannot occur. We show that there must be a preference cycle in $\succsim_N \in \mathcal{R}^X$ contrary to the inductive assumption. Hence, to prevent the 3-cycle, one of three sub-cases must occur:

- (a) $Y_{1,2} \prec_1 Y$ – using (2.6) we have that $Y_{m,1} \succsim_1 Y_{1,2} \prec_1 Y$. Since Y is in the m -cycle, we obtain a cycle, contrary to the inductive assumption.
- (b) $Y \prec_3 Y_{2,3}$ – using (2.6) we have that $Y \prec_3 Y_{2,3} \succsim_3 Y_{3,4}$. Since Y is in the m -cycle, we obtain a cycle, contrary to the inductive assumption.
- (c) $Y \sim_1 Y_{1,2} \sim_2 Y_{2,3} \sim_3 Y$ – using a strict preference from (2.6), we obtain a cycle, contrary to the inductive assumption.

2. $Y \neq Y_{l,l+1}$ for any $l = 1, \dots, m$. We show that there must be a preference cycle in another profile $\succsim'_N \in \mathcal{R}^X$ contrary to the inductive assumption. Since contracts

are heterogeneous, there must exist a preference profile $\succsim'_N \in \mathcal{R}^X$ such that:

$$Y_{m,1} \succsim'_1 Y \succsim'_1 Y_{1,2} \quad (2.7)$$

(Y has a contract between i_1 and i_3) and since Y is not in the cycle, all the \succsim'_N -preferences of agents in the m -cycle are the same as \succsim_N -preferences. Since all the preferences for agents outside Y remain the same under \succsim'_N , it must be the case that

$$Y \succsim'_1 Y_{1,2} \succsim'_2 Y_{2,3} \quad (2.8)$$

There are two possible cases that must occur:

- (a) $Y \prec'_3 Y_{2,3}$ – using (2.6) under \succsim'_N -preferences, we have $Y \prec'_3 Y_{2,3} \succsim'_4 Y_{3,4}$. Now using (2.7), we obtain $Y \prec'_3 Y_{2,3} \succsim'_4 Y_{3,4} \succsim' \dots \succsim' Y_{m,1} \succsim'_1 Y$ contrary to the inductive assumption.
- (b) $Y_{2,3} \succsim'_3 Y$ – using (2.6) under \succsim'_N -preferences, we have $Y \succsim'_1 Y_{1,2} \succsim'_2 Y_{2,3} \succsim'_3 Y$. Hence, we obtain $Y \succsim'_3 Y_{2,3} \succsim'_4 Y_{3,4} \succsim' \dots \succsim' Y_{m,1} \succsim'_1 Y$. But since all the \succsim'_N -preferences of agents in the m -cycle in (2.6) are the same as \succsim_N -preferences, there is a cycle under \succsim'_N -preferences, contrary to the inductive hypothesis.

Hence, there is a contradiction in every case, which completes the proof. □

Proof of Theorem 2.2. Follows immediately from Lemmata 2.3, 2.4, and 2.2. □

2.7.2 Necessity of strong pairwise alignment for a strongly stable contract allocation

Lemma 2.5. *If contract set X is very heterogeneous and all preference profiles in \mathcal{R}^X admit a strongly stable contract allocation, then there is no preference profile $\succsim_N \in \mathcal{R}^X$*

which contains a 3-cycle $Y_{1,3} \prec_1 Y_{1,2} \succsim_2 Y_{2,3} \succsim_3 Y_{1,3}$.

Proof. This lemma does not require strong pairwise alignment to hold.

Towards contradiction, assume that there exists a preference profile which admits a 3-cycle in the domain. We aim to show that for some other preference profile in the domain in which there is no strongly stable contract allocation.

First, we show that there is a preference profile with a strict 3-cycle for agents i_1, i_2 and i_3 and then we show that there is a preference profile such that for any Y which contains agents who share contracts with $l = 1, 2, 3$ in $Y_{l,l+1}$ (l modulo 3), all these agents strictly prefer one of the contract allocations in the strict 3-cycle to Y .

If all three contract allocations $Y_{1,2}, Y_{2,3}, Y_{3,1}$ are the same for i_1, i_2 and i_3 , then their preferences are cyclic. Suppose these contract allocations are different. Since contracts are very heterogeneous, a new preference profile \succsim'_N can be found so that:

$$Y_{3,1} \prec'_1 Y_{1,2}, Y_{1,2} \prec'_2 Y_{2,3}, Y_{2,3} \succ'_3 Y_{3,1} \quad (2.9)$$

because the \succsim'_N -preferences of agents between contract allocations that do not include $Y_{1,2}$ must remain the same as \succsim_N . Now repeat this on \succsim'_N to obtain:

$$Y_{3,1} \prec''_1 Y_{1,2}, Y_{1,2} \prec''_2 Y_{2,3}, Y_{2,3} \prec''_3 Y_{3,1} \quad (2.10)$$

Since contracts are very heterogeneous, for any agent i who has a contract with either i_1 or i_2 in $Y_{1,2}$, or with either i_2 or i_3 in $Y_{2,3}$, or with either i_1 or i_3 in $Y_{1,3}$ and for any $Y_i \neq Y_{1,2}, Y_{2,3}, Y_{3,1}$, we can find a preference profile $\succsim_N^* \in \mathcal{R}^X$ such that $Y_i \prec_i^* Y_{l,l+1}$ whenever $i \in a(Y) \cap a(Y_{l,l+1})$ for $l = 1, 2, 3$ (l modulo 3).

The resulting preference profile $\succsim_N^* \in \mathcal{R}^X$ does not admit a strongly stable allocation because any Y is blocked, which contradicts the initial assumption. \square

Lemma 2.6. *Suppose that contract set is X heterogeneous and no preference profile in \mathcal{R}^X admits a 3-cycle from Lemma 2.5. Then all preference profiles in \mathcal{R}^X satisfy strong pairwise alignment.*

Proof. We simply need to show that for any contract allocations $Y \neq Y'$ such that $Y = Y_{i,j}$ and $Y' = Y'_{i,j}$ if $Y \succsim_i Y'$, then for any $a(Y_i) \cap a(Y'_i) = j$, $Y \succsim_j Y'$. We can assume that not every agent is a member of every contract in either Y or Y' (i.e. $Y, Y' \notin \mathcal{Y}$).

By contractual feasibility, there exists a contract allocation Y'' such that *only* i and k share contracts. Since contracts are heterogeneous, we pick any preference profile such that:

$$Y \succsim_i Y'' \succsim_i Y' \quad (2.11)$$

Towards a contradiction, assume $Y \succ_j Y'$. Then, by contractual feasibility and heterogeneity of the contract set, there exists Y''' , in which only j and k share contracts:

$$Y' \succsim_j Y''' \succsim_j Y \quad (2.12)$$

If $Y' \prec_j Y'''$, and $Y'' \prec_k Y'''$, then we get a prohibited 3-cycle using (2.11) and (2.12):

$$Y \succsim_i Y'' \prec_k Y''' \succsim_j Y \quad (2.13)$$

If $Y' \prec_j Y'''$, and $Y''' \succsim_k Y''$, then we get a prohibited 3-cycle using (2.11) and (2.12):

$$Y' \prec_j Y''' \succsim_k Y'' \succsim_i Y' \quad (2.14)$$

If $Y''' \prec_j Y$, and $Y'' \succ_k Y'''$, then we get a prohibited 3-cycle using (2.11):

$$Y_j \succ_i Y'' \succ_k Y''' \prec_j Y \quad (2.15)$$

Finally, if $Y''' \prec_j Y$, and $Y''' \prec_k Y''$, then we get a prohibited 3-cycle using (2.11) and (2.12):

$$Y' \succ_j Y''' \prec_k Y'' \succ_i Y' \quad (2.16)$$

In every case, there is contradiction and this completes the proof. □

Proof of Theorem 2.3. Follows immediately from Lemmata 2.5 and 2.6. □

References

- Alkan, A. (1988). Nonexistence of stable threesome matchings. *Mathematical Social Sciences* 16(2), 207 – 209.
- Aumann, R. J. (1959). Acceptable points in general cooperative n-person games. In *Contributions to the Theory of Games IV*, Annals of Mathematics Study 40, pp. 287–324. Princeton University Press, Princeton.
- Baccara, M., A. İmrohoroğlu, A. J. Wilson, and L. Yariv (2012). A field study on matching with network externalities. *American Economic Review* 102(5), 1773–1804.
- Baker, G. P., R. Gibbons, and K. J. Murphy (2008). Strategic alliances: Bridges between “Islands of Conscious Power”. *Journal of the Japanese and International Economies* 22(2), 146–163.
- Baldwin, E. and P. Klemperer (2013). Tropical geometry to analyse demand. Discussion paper, University of Oxford, Department of Economics.
- Bando, K. (2012). Many-to-one matching markets with externalities among firms. *Journal of Mathematical Economics* 48(1), 14 – 20.
- Blair, C. (1988). The lattice structure of the set of stable matchings with multiple partners. *Mathematics of Operations Research* 13(4), 619–628.
- Crawford, V. P. and E. M. Knoer (1981). Job matching with heterogeneous firms and workers. *Econometrica* 49(2), 437–450.
- Demange, G. (2004). On group stability in hierarchies and networks. *Journal of Political Economy* 112(4), 754–778.
- Dimitrov, D. and E. Lazarova (2011). Two-sided coalitional matchings. *Mathematical Social Sciences* 62(1), 46–54.
- Dutta, B. and J. Massó (1997). Stability of matchings when individuals have preferences over colleagues. *Journal of Economic Theory* 75(2), 464–475.
- Dutta, B. and S. Mutuswami (1997). Stable networks. *Journal of Economic Theory* 76(2), 322–344.
- Echenique, F. (2012). Contracts vs. salaries in matching. *American Economic Review* 102(1), 594–601.
- Echenique, F. and J. Oviedo (2004). Core many-to-one matchings by fixed-point methods. *Journal of Economic Theory* 115(2), 358–376.

- Echenique, F. and J. Oviedo (2006). A theory of stability in many-to-many matching markets. *Theoretical Economics* 1(1), 233–273.
- Echenique, F. and M. B. Yenmez (2007). A solution to matching with preferences over colleagues. *Games and Economic Behavior* 59(1), 46–71.
- Erdil, A. and T. Kumano (2011). Stability and efficiency in the general priority-based assignment. Technical report, Mimeo.
- Flanagan, F. (2013). Contracts v. preferences over colleagues in matching. Mimeo, University of Wisconsin-Madison.
- Fox, J. T. (2009, October). Estimating matching games with transfers. Working Paper 14382, NBER.
- Gale, D. and L. S. Shapley (1962). College admissions and the stability of marriage. *American Mathematical Monthly* 69(1), 9–15.
- Goyal, S. (2009). *Connections: An Introduction to the Economics of Networks*. Princeton University Press.
- Grout, P. (1984). Investment and wages in the absence of binding contracts: A Nash bargaining approach. *Econometrica* 52(2), 449–460.
- Hatfield, J. W. and F. Kojima (2010). Substitutes and stability for matching with contracts. *Journal of Economic Theory* 145(5), 1704–1723.
- Hatfield, J. W. and S. D. Kominers (2011a). Contract design and stability in matching markets. Working paper, Harvard Business School.
- Hatfield, J. W. and S. D. Kominers (2011b). Multilateral matching. Mimeo, Stanford GSB.
- Hatfield, J. W. and S. D. Kominers (2012). Matching in networks with bilateral contracts. *American Economic Journal: Microeconomics* 4(1), 176–208.
- Hatfield, J. W. and P. Milgrom (2005). Matching with contracts. *American Economic Review* 95(4), 913–935.
- Hellman, T. (1998). The allocation of control rights in venture capital contracts. *RAND Journal of Economics* 29(1), 57–76.
- Jackson, M. O. (2005). Allocation rules for network games. *Games and Economic Behavior* 51(1), 128–154.
- Jackson, M. O. (2008). *Social and Economic Networks*. Princeton University Press.
- Jackson, M. O. and A. van den Nouweland (2005). Strongly stable networks. *Games and Economic Behavior* 51(2), 420–444.

- Jackson, M. O. and A. Watts (2010). Social games: Matching and the play of finitely repeated games. *Games and Economic Behavior* 70(1), 170–191.
- Jackson, M. O. and A. Wolinsky (1996). A strategic model of social and economic networks. *Journal of Economic Theory* 71(1), 44–74.
- Kalai, E. and M. Smorodinsky (1975). Other solutions to Nash’s bargaining problem. *Econometrica* 43(3), 513–518.
- Kelso, A. S. and V. P. Crawford (1982). Job matching, coalition formation, and gross substitutes. *Econometrica* 50(6), 1483–1504.
- Klaus, B. and M. Walzl (2009). Stable many-to-many matchings with contracts. *Journal of Mathematical Economics* 45(7-8), 422–434.
- Kojima, F. and M. U. Ünver (2008). Random paths to pairwise stability in many-to-many matching problems: a study on market equilibration. *International Journal of Game Theory* 36(3-4), 473–488.
- Kominers, S. D. (2010). Matching with preferences over colleagues solves classical matching. *Games and Economic Behavior* 68(2), 773–780.
- Kominers, S. D. (2012). On the correspondence of contracts to salaries in (many-to-many) matching. *Games and Economic Behavior* 75(2), 984 – 989.
- Kongo, T. (2011). Value of games with two-layered hypergraphs. *Mathematical Social Sciences* 62(2), 114–119.
- Konishi, H. and M. U. Ünver (2006). Credible group stability in many-to-many matching problems. *Journal of Economic Theory* 129(1), 57–80.
- Macho-Stadler, I., D. Pérez-Castrillo, and R. Veugelers (2008). Designing contracts for university spin-offs. *Journal of Economics & Management Strategy* 17(1), 185–218.
- Myerson, R. B. (1977). Graphs and cooperation in games. *Mathematics of Operations Research* 2(3), 225–229.
- Myerson, R. B. (1980). Conference structures and fair allocation rules. *International Journal of Game Theory* 9(3), 169–182.
- Newman, M. E. J. (2010). *Networks: An Introduction*. Oxford University Press: Oxford.
- Ostrovsky, M. (2008). Stability in supply chain networks. *American Economic Review* 98(3), 897–923.
- Oswald, A. (1982). The microeconomic theory of the trade union. *Economic Journal* 92(367), 576–595.

- Page Jr., F. H. and M. Wooders (2010). Club networks with multiple memberships and noncooperative stability. *Games and Economic Behavior* 70(1), 12–20.
- Pycia, M. (2012). Stability and preference alignment in matching and coalition formation. *Econometrica* 80(1), 323–362.
- Revilla, P. (2007). Many-to-one matching when colleagues matter. Working paper 146, Fondazione Eni Enrico Mattei.
- Roth, A. E. (1984b). Stability and polarization of interests in job matching. *Econometrica* 52(1), 47–58.
- Roth, A. E. (1991). A natural experiment in the organization of entry-level labor markets: Regional markets for new physicians and surgeons in the United Kingdom. *American Economic Review* 81(3), 415–440.
- Roth, A. E. and M. Sotomayor (1990). *Two-sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Cambridge University Press.
- Sasaki, H. and M. Toda (1996). Two-sided matching problems with externalities. *Journal of Economic Theory* 70(1), 93–108.
- Shapley, L. S. (1967). On balanced sets and cores. *Naval Research Logistics Quarterly* 14(4), 453–460.
- Shapley, L. S. (1971). Cores of convex games. *International Journal of Game Theory* 1(1), 11–26.
- Slikker, M., B. Dutta, A. van den Nouweland, and S. Tijs (2000). Potential maximizers and network formation. *Mathematical Social Sciences* 39(1), 55–70.
- Smith, A. (1970, 1776). *An Inquiry into the Nature and Causes of The Wealth of Nations*. Penguin.
- Sotomayor, M. (1999). Three remarks on the many-to-many stable matching problem. *Mathematical Social Sciences* 38, 55–70.
- Sun, N. and Z. Yang (2006). Equilibria and indivisibilities: Gross substitutes and complements. *Econometrica* 74(5), 1385–1402.

Chapter 3

Gross substitutes and complements: a generalisation

Abstract

This chapter extends the gross substitutes and complements (GSC) framework of Sun and Yang (2006) to a more general substitutes and complements structure. We show that competitive equilibrium exists under significantly weaker, easily checkable, and interpretable conditions. We also consider the supply-chain structure of Ostrovsky (2008) and show that chain-stable networks may exist even in the presence of contractual cycles provided the market is appropriately segmented.

Keywords: competitive equilibrium, substitutes, complements, chain stability, networks, contracts.

JEL Classification: *D51, C7, L14*

3.1 Introduction

As economies grow, compete, and specialise, they exhibit more complements in production and more substitutes in consumption.¹ Moreover, many markets exhibit natural discreteness and heterogeneity. Car manufacturers match with hundreds of suppliers which deliver unique engine and suspension parts for different models. Highly-skilled workers are often headhunted by firms in niche industries with thin labour markets. In designed markets, spectrum rights are sold in separate bands, and different bidders may be interested in one or more units. However, it is well known that, in markets with indivisible commodities, competitive equilibrium does not always exist when complementarities are present (Kelso and Crawford, 1982, Gul and Stacchetti, 1999). This chapter offers a new sufficient condition for the existence of competitive equilibrium in an endowment economy in which agents trade indivisible substitute and complement goods. In our model, all goods can be partitioned into sets of substitutes and every buyer regards goods from some two partition elements as complements. As an example, consider an economy in which the seller offers three types of goods: jackets, trousers, and shirts. Buyers view any type of good as substitutes – this is a natural assumption when the goods of a particular type are sufficiently similar. There are also two types of buyer: a student who views jackets and trousers as complements, and a professor who views trousers and shirts as complements. We show that in this sort of economy competitive equilibrium is guaranteed to exist. However, if we add another agent into the economy – a post-doc who regards jackets and shirts as complements – then competitive equilibrium is no longer guaranteed to exist (see section 3.3.4.2). This failure of equilibrium existence occurs because there is an *odd cycle* in the *substitutes and complements (SC) structure* of agents' preferences: jackets and trousers are complements

¹Milgrom and Roberts (1990) gives numerous examples of how manufacturers expand their production lines and increasingly rely on independent suppliers.

for the student, trousers and shirts are complements for the professor, and shirts and jackets are complements for the post-doc. In section 3.3, we show that competitive equilibrium exists whenever these odd cycles are absent: a much weaker, yet intuitive, condition than those found previously.²

In section 3.4, we study matching markets in which firms have preferences over contracts with other firms. In this model, contract terms are drawn from a finite set.³ A matching market with contracts is similar to a trading economy with indivisible goods, however, in the contracts model, we drop the assumption of continuous prices. We consider *contractual networks* in which firms form downstream contracts to sell outputs and form upstream contracts to buy inputs (Ostrovsky, 2008). In some markets, such as electricity markets (Hatfield and Kominers, 2012) and industries with complex supply chains, firms simultaneously supply inputs to *and* buy outputs from other firms (possibly through intermediaries). We say that contractual networks in these markets exhibit *contractual cycles*. In the absence of continuous prices, we depart from the competitive equilibrium notion. Instead, we show that *chain-stable* (pairwise stable) *contractual networks* always exist as long as the no-odd-cycles-condition is satisfied over contractual cycles as long as the market is appropriately segmented. There is a natural analogy to the trading economy because we assume that firms regard their upstream and downstream contracts (which form sets of substitutes) as complements. Our unified approach reveals interesting connections between equilibria in trading economies with divisible money and stable allocations in matching markets with contracts.

²The “no odd party” (Tan, 1991) and “no-odd-rings” conditions (Chung, 2000) guarantee existence of stable matchings in the roommate market. Gudmundsson (2013) showed that absence of odd cycles in a certain linear programming problem guarantees existence of equilibrium in the partnership formation problem (Talman and Yang, 2011). However, these results for one-sided matching problems are logically unrelated to the present chapter.

³See a further discussion of this in Chapter 2 of this thesis.

3.2 Related literature

This chapter is inspired by and extends the results of the *gross substitutes and complements* (GSC) preference framework (Sun and Yang, 2006). The authors showed that if, for example, a seller offers trousers and shirts and all buyers regard any two shirts (or any two pairs of trousers) as substitutes, but any shirt and pair of trousers as complements, then competitive equilibrium will exist in this economy when agents' utility functions are quasilinear in prices. In the case of a trading economy, Sun and Yang (2006) is a special case of our model where there are two types of goods viewed as substitutes by buyers (e.g. *only* trousers and shirts), and the seminal contribution of Kelso and Crawford (1982) is a special case where there is only one type of substitute good (e.g. trousers). Other generalisations of the GSC framework were proposed by Baldwin and Klemperer (2013) and Shioura and Yang (2013).

Ostrovsky (2008) developed the intuition of the GSC framework to study *supply chains* in which firms sign finite contracts (Crawford and Knoer, 1981, Roth, 1984b, Hatfield and Milgrom, 2005). In a supply chain, there are agents, who only supply inputs (e.g. farmers); agents, who only buy final outputs (e.g. consumers); while the rest of the agents are intermediaries, who buy inputs and sell outputs (e.g. supermarkets). We extend this supply-chain framework in which agents' preferences satisfy *same-side substitutability* and *cross-side complementarity* (also jointly called *full substitutability* by Hatfield and Kominers, 2012) to the case of more general contractual networks with cycles. Hatfield and Kominers (2012) showed that if a contractual network has a contractual cycle – the existence of an agent, who buys inputs from and sells outputs to another agent (perhaps via intermediaries) – then a *stable* contract allocation may fail to exist. We show that, in the case of *even* contractual cycles, allocations that satisfy a weaker notion of stability – *chain* (pairwise) *stability* – can still be found provided the

market structure is appropriately segmented. We can conjecture that by strengthening assumptions about the preferences over contracts, a chain-stable outcome can be found in more general contractual networks.

3.3 Model of a trading economy

First, we consider a standard trading economy in which agents are endowed with utility functions quasilinear in prices. The relationship between this trading economy, auctions, and matching markets has been known in the literature for some time (Kelso and Crawford, 1982, Milgrom, 2000, Sun and Yang, 2006, Milgrom and Strulovici, 2009). More recently, Hatfield et al. (2013) and Drexler (2013) used the insights of this framework to model trading networks with continuous prices. We can apply the substitutes and complements (SC) structure, described below, to generalise all these models (see section 3.6.1).

3.3.1 Ingredients

There is a finite set of agents $i \in I$ and a finite set of indivisible goods $\omega \in \Omega$ in the economy. Goods are partitioned into M (possibly empty) disjoint subsets (of similar goods) of Ω , forming a set $\mathcal{S} = \{S_1, \dots, S_M\}$ such that $S_n \cap S_m = \emptyset$ (where $n, m = 1, \dots, M; n \neq m$) and $\bigcup_{m=1}^M S_m = \Omega$. Each element of the partition represents a set of similar goods (such as shirts of a different colour). Let $\Psi \in 2^\Omega$ be a bundle of goods and Ψ_i be a bundle for agent $i \in I$. Denote p_ω as the price of good ω and $p \in \mathbb{R}^{|\Omega|}$ as the price vector. An allocation is a partition Π of goods into (possibly empty) bundles for different agents ($\Pi = \{\Psi_i\}_{i \in I}$ such that $\bigcup_{i \in I} \Psi_i = \Omega$ and $\Psi_i \cap \Psi_j = \emptyset$). An arrangement is a pair $[\Pi, p]$, which associates prices to goods in the economy. In this

section, we assume that agents' utility functions are quasilinear in prices

$$U_i([\Pi; p]) \equiv u_i(\Psi_i) - \sum_{\omega \in \Psi_i} p_\omega \quad (3.1)$$

and agents are not subject to any liquidity or budget constraints. Therefore, the trading economy can be described by the set of goods and the agents' valuations of every bundle: $\mathcal{E} = \{\Omega, (u_i, i \in I)\}$. The demand correspondence $D_i : \mathbb{R}^{|\Omega|} \rightarrow 2^\Omega$ for agent i is defined in the usual way

$$D_i(p) \equiv \arg \max_{\Psi \subseteq \Omega} U_i([\Psi; p]) \quad (3.2)$$

A competitive equilibrium in this economy consists of an allocation of goods to agents and a vector of prices, such that the market clears: every agent demands precisely his allocation at this price vector.

Definition 3.1. *Competitive equilibrium* is an arrangement $[\Pi; p]$ such that for all $i \in I$, $\Psi_i \in D_i(p)$.

3.3.2 Preferences

First, we formally define preferences that satisfy gross substitutes and complements (Sun and Yang, 2006). Let us consider $\mathcal{S} = \mathcal{S}^* = \{S_1, S_2\}$, which represents shirts (S_1) and trousers (S_2). For notation purposes, define $e(k)$ is a k th unit vector in $\mathbb{R}^{|\Omega|}$ and $A^c = \Omega \setminus A$.

Definition 3.2 (Sun and Yang, 2006). Preferences of agent i satisfy *gross substitutes and complements (GSC)* on \mathcal{S}^* if for any $p \in \mathbb{R}^{|\Omega|}$, $\omega_k \in S_m$, $\delta \geq 0$, $A \in D_i(p)$, there exists $B \in D_i(p + \delta e(k))$ such that $[A \cap S_m] \setminus \{\omega_k\} \subseteq B$ and $A^c \cap S_m^c \subseteq B^c$.

In other words, preferences of agent i satisfy GSC if whenever the price of a shirt (trousers) increases, i 's resulting demand, B , for other shirts (trousers) does not fall i.e

$[A \cap S_m] \setminus \{\omega_k\} \subseteq B$, and demand for trousers (shirts) does not rise i.e. $A^c \cap S_m^c \subseteq B^c$. Using our example in the Introduction, the professor has GSC preferences over shirts and trousers. Sun and Yang (2006) show that whenever preferences of all agents satisfy GSC competitive equilibrium exists. We now turn to the main assumption on individual preferences and structure of the economy which generalises GSC.

Definition 3.3. Preferences have a *substitutes and complements (SC) structure* on \mathcal{S} if for any $p \in \mathbb{R}^{|\Omega|}$, $\omega_k \in S_m$, $\delta \geq 0$, $A \in D_i(p)$, $S_m \in \mathcal{S}$ and $i \in I$, there exists one $S_n \in \mathcal{S}$ and $B \in D_i(p + \delta e(k))$ such that $[A \cap S_m] \setminus \{\omega_k\} \subseteq B$, $A^c \cap S_n \subseteq B^c$ and $A \cap [S_m \cup S_n]^c = B \cap [S_m \cup S_n]^c$.

In words, agents' demand correspondences have a SC structure if we can divide goods into a partition \mathcal{S} (for all agents) such that, whenever we consider preferences over goods contained in any two elements of \mathcal{S} *in isolation*, these preferences satisfy GSC for some agents. From now on, whenever we say that agents have GSC preferences over S_m and S_n in the SC structure, we will mean that these agents would have GSC preferences over $\mathcal{S}^* = \{S_m, S_n\}$ if the goods in S_m and S_n were considered in isolation. Different agents may have GSC preferences over different pairs of the partition elements of \mathcal{S} .⁴ Again returning to our example, in a SC structure, the student has GSC preferences over jackets and trousers and the post-doc has GSC preferences over jackets and shirts. It is worth noting that, for agents who have GSC preferences over S_m and S_n , changes in prices for a good contained in S_m do not have any effect on the demands for goods outside $S_n \cup S_m$, and changes in prices for goods contained $[S_n \cup S_m]^c$ do not affect the demands for goods in $S_n \cup S_m$. In other words, for these agents, goods in $S_n \cup S_m$ and $[S_n \cup S_m]^c$ are independent. To see how the SC structure generalises GSC, note that preferences with a SC structure satisfy GSC for all agents if $M = 2$. Since we do not

⁴We could allow the same agent to have GSC preferences over several pairs of elements of \mathcal{S} (because of the quasilinearity of the utility functions), but, while this complicates the exposition, it does not affect the results in any way.

rule out that $S = \emptyset$, it is clear that if $M = 2$ and $\mathcal{S} = \{S, \emptyset\}$, we return to the *gross substitutes* framework of Kelso and Crawford (1982).⁵

3.3.3 SC cycles

We now introduce a formal definition of an odd SC cycle, which was described in the Introduction.

Definition 3.4. Preferences form a *SC cycle* on \mathcal{S} if

1. Preferences have a *SC structure* on \mathcal{S} , and
2. Non-empty elements of \mathcal{S} can be arranged in some order $S_1 \dots S_K$ ($2 \leq K \leq M$) such that there exists an agent with GSC preferences over goods contained in S_m and S_{m+1} for $m = 1, \dots, K$ ($m \pmod K$).

Preferences may form SC cycles of different lengths. In Sun and Yang’s framework, the largest SC cycle is of length 2. We want to focus on the existence of SC cycles of odd length.

Definition 3.5. Preferences form an *odd SC cycle* on \mathcal{S} if preferences form a *SC cycle* on \mathcal{S} and there is an odd $K > 2$.

Hence, preferences of the professor, the student, and the post-doc over the clothes which we described in the Introduction form an odd SC cycle with $M = K = 3$.

3.3.4 Motivating examples

Will a competitive equilibrium exist in a trading economy where agents’ preferences have a SC structure? We present two examples to illustrate that in general the answer is “no”.

Table 3.1: Competitive equilibrium with no odd cycles in the SC structure

Ψ	$u_i(\Psi_i)$	$u_j(\Psi_j)$	$u_k(\Psi_k)$	$u_l(\Psi_l)$
\emptyset and otherwise	0	0	0	0
$\{\omega_1, \omega_2\}$	1	0	0	0
$\{\omega_2, \omega_3\}$	0	1	0	0
$\{\omega_3, \omega_4\}$	0	0	1	0
$\{\omega_4, \omega_1\}$	0	0	0	1
$\{\omega_1, \omega_2, \omega_3\}$	1	1	0	0
$\{\omega_2, \omega_3, \omega_4\}$	0	1	1	0
$\{\omega_1, \omega_3, \omega_4\}$	0	0	1	1
$\{\omega_1, \omega_2, \omega_4\}$	1	0	0	1
$\{\omega_1, \omega_2, \omega_3, \omega_4\}$	1	1	1	1

3.3.4.1 Equilibrium with no odd cycles

Consider a trading economy with four buyers i, j, k, l , a seller s , and four goods $\{\omega_1, \omega_2, \omega_3, \omega_4\}$. The seller's values are zero for every bundle. Buyers get utility 1 from getting a bundle containing their two desired goods (the bundle may include other goods) and zero otherwise as illustrated in Table 3.1. Agent i has GSC preferences over $\{\omega_1, \omega_2\}$; j over $\{\omega_2, \omega_3\}$; k over $\{\omega_3, \omega_4\}$; and l over $\{\omega_4, \omega_1\}$. Hence, $S_m = \{\omega_m\}$ for $m = 1, \dots, 4$ and preferences satisfy a SC structure. Suppose that the price vector is $p^* = (p_1, p_2, p_3, p_4) = (1, 0, 1, 0)$ and $\Pi^* = \{\Psi_i, \Psi_j, \Psi_k, \Psi_l\} = \{\{\omega_1, \omega_2\}, \emptyset, \{\omega_3, \omega_4\}, \emptyset\}$. Bundle $\{\omega_1, \omega_2\}$ is allocated to agent i and k gets bundle $\{\omega_3, \omega_4\}$; agents j and l get nothing. Arrangement $[\Pi^*; p^*]$ constitutes a competitive equilibrium.

It is easy to show that a slight variation of the double-track auction proposed by Sun and Yang (2008, 2009) can converge to this price vector. Suppose, a seller announces an initial price vector $p(0) = (2, 0, 2, 0)$ i.e. selecting high prices for ω_1 and ω_3 and low

⁵Therefore, we do not rule out that some agents may only view goods within an element \mathcal{S} as substitutes and no goods as complements.

Table 3.2: No competitive equilibrium with an odd cycle in the SC structure

Ψ	$u_i(\Psi_i)$	$u_j(\Psi_j)$	$u_k(\Psi_k)$
\emptyset and otherwise	0	0	0
$\{\omega_1, \omega_2\}$	1	0	0
$\{\omega_2, \omega_3\}$	0	1	0
$\{\omega_3, \omega_1\}$	0	0	1
$\{\omega_1, \omega_2, \omega_3\}$	1	1	1

prices for ω_2 and ω_4 . Now, ω_1 and ω_3 are under-demanded and ω_2 and ω_4 are balanced (according to the auction rules). Therefore, the seller reduces the prices of ω_1 and ω_3 and keeps the prices of ω_2 and ω_4 constant. At $p(1) = (1, 0, 1, 0)$, all goods are balanced, so this price vector constitutes a competitive equilibrium.

3.3.4.2 No competitive equilibrium with odd cycles

This example is adapted from Bikhchandani and Mamer (1997) and illustrates that if preferences contain odd cycles in the SC structure no competitive equilibrium may exist in a trading economy.

Consider three buyers i, j, k , a seller s , and three goods $\omega_1, \omega_2, \omega_3$. The seller's values are zero for every bundle. Agents get utility 1 from getting a bundle containing their two desired goods (the bundle may include other goods) and zero otherwise as illustrated in Table 3.2. Once again, $S_m = \{\omega_m\}$ for $m = 1, \dots, 3$ and preferences satisfy a SC structure. The problem is symmetric so without loss of generality assume that in a competitive equilibrium allocation $\Psi_i = \{\omega_1, \omega_2, \omega_3\}$.⁶ Then, it must be the

⁶We only need to consider efficient allocations.

case that

$$p_1 + p_2 + p_3 \leq 1 \tag{3.3}$$

$$p_2 + p_3 \geq 1 \tag{3.4}$$

$$p_3 + p_1 \geq 1 \tag{3.5}$$

otherwise agent i would not demand $\{\omega_1, \omega_2, \omega_3\}$ and agents j or k would demand a bundle. However, adding equations (3.4) and (3.5) and using equation (3.3), we obtain

$$p_1 + p_2 + 2p_3 \geq 2 \tag{3.6}$$

$$2 - 2p_3 \leq p_1 + p_2 \leq 1 - p_3 \tag{3.7}$$

which only holds when $p_3 = 1$ implying that $p_1 + p_2 \leq 0$. Hence, the seller would not sell the bundle and this allocation cannot be supported by market clearing prices.⁷

3.4 Main result

We now present the first main result of this chapter.⁸

Theorem 3.1. *Suppose that in the trading economy $\mathcal{E} = \{\Omega, (u_i, i \in I)\}$ agents' preferences have a SC structure on \mathcal{S} and do not form any odd SC cycles. Then this trading economy has a competitive equilibrium.*

⁷Another candidate is for a competitive equilibrium allocation is, without loss of generality, $\Psi_i = \{\omega_1, \omega_2\}$, leaving ω_3 unsold. This implies $p_3 = 0$. However, the following inequalities must hold

$$\begin{aligned} p_1 + p_2 &\leq 1 \\ p_2 + p_3 &\geq 1 \\ p_3 + p_1 &\geq 1 \end{aligned}$$

but this implies that $p_1 \geq 1$, $p_2 \geq 1$, and $2 \leq p_1 + p_2 \leq 1$, which means that this allocation cannot be supported in a competitive equilibrium either.

⁸This result was simultaneously obtained by Baldwin and Klemperer (2013) and Sun and Yang (2011).

Proof. Take any element in $\mathcal{S} = (S_1, \dots, S_M)$. Consider a graph G , whose nodes are the elements of \mathcal{S} and whose edges represent the fact that some agents have GSC preferences between these two nodes. Since \mathcal{S} does not form an odd cycle, it is a bipartite graph whose nodes can be divided into two disjoint sets S^A and S^B such that every edge connects a node in S^A to the node(s) in S^B only. Consider any agent $i \in I$. Without loss of generality, assume agent i has GSC preferences over S_m and S_{m+1} and that S_m is a node in S^A (call it S_m^A) and S_{m+1} is a node in S^B (call it S_{m+1}^B for symmetry). Consider the valuation function $v_i(T) = u_i(T \cap (S_m^A \cup S_{m+1}^B))$ for any goods in $T \subseteq S^A \cup S^B$. Agent i 's preferences described by utility function $V_i(T)$ satisfy GSC over S^A and S^B . Hence, the utility function of every $i \in I$ satisfies GSC over S^A and S^B . Therefore, following Sun and Yang (2006) the general trading economy on graph G must have at least one competitive equilibrium. \square

We can illustrate the proof and the examples above using graphs. In Figures 3.1 and 3.2, nodes represent elements of \mathcal{S} containing substitute goods. Edges represent the fact that some agents regard goods from two elements of \mathcal{S} as complements. In Figure 3.1, there is an odd SC cycle of size 3 (described in the Introduction) and competitive equilibrium may not exist. In Figure 3.2, there are only even cycles and therefore the existence of competitive equilibrium is guaranteed for any valuations that satisfy that SC structure. It is worth emphasising that we do not demand acyclicity of the SC structure. A nonempty connected graph is acyclic if and only if the number of edges is exactly one less than the number of nodes. A simple corollary of Theorem 3.1 is that any trading economy in which preferences have an acyclic SC structure has a competitive equilibrium.

Figure 3.1: Odd SC cycle





Figure 3.2: Even SC cycle



Trading economy: nodes represent non-empty elements of \mathcal{S} ; edges represent GSC preferences between these elements.

3.5 Application: contractual networks

We apply the insights of Theorem 3.1 to the case of supply chains, which were introduced by Ostrovsky (2008) and were more recently studied by Westkamp (2010) and Hatfield and Kominers (2012). First, we drop the assumption of quasilinearity and continuous prices and consider fully substitutable preferences over finite contracts (which are an analogue of gross substitutes and complements). Second, we depart from the competitive equilibrium notion and follow Ostrovsky (2008) in considering chain-stable outcomes.

3.5.1 Model

There is finite set of firms I and a finite set of contracts Ω . Firms are exogenously partitioned into M non-empty *segments* $\mathcal{C} = \{C_1, \dots, C_M\}$ (so $\cup_{j=1}^M C_j = I$). A contract $\omega \in \Omega$ is an agreement between a buyer $\omega_b \in I$ and a seller $\omega_s \in I$. Hence, $\omega_I \equiv \{\omega_b, \omega_s\}$ is the set of firms associated with contract ω and, more generally, Ω'_I is the set of firms associated with contract set $\Omega' \subseteq \Omega$. Call $B_i = \{\omega \in \Omega \mid \omega_b = i\}$ and

$S_i = \{\omega \in \Omega | \omega_s = i\}$ the sets of i 's upstream and downstream contracts – for which i is a buyer and a seller, respectively. Clearly, B_i and S_i form a partition over the set of contracts $\Omega_i \equiv \{\omega \in \Omega | i \in \omega_I, i \in I\}$ which involve i , since an agent cannot be a buyer and a seller in the same contract.


We now impose some conditions on the market structure or, more specifically, on Ω . First, for all $\omega \in \Omega$ if $\omega_b \in C_j$ then $\omega_s \notin C_j$, i.e. no two firms in the same segment may form a contract. Second, for all $j = 1, \dots, M$, all firms in segment C_j may only form downstream contracts with firms in the downstream segment C_{j+1} and upstream contracts only with firms in the upstream segment C_{j-1} (where $C_{M+1} \equiv C_1$ and $C_0 \equiv C_M$). This market structure differs from the supply chain structure investigated by Ostrovsky (2008) in two ways. Our structure does not rule out the possibility of a firm being an indirect supplier to an upstream firm. However, we only allow firms to form contracts in the downstream or the upstream segments. In particular, we do not allow a firm to be a direct **and** an indirect supplier (buyer) to (from) a downstream (upstream) firm. We illustrate this prohibited structure in Figure 3.3, where there are three segments in the market. We label each node as an agent and each directed edge (arc) as a contract between them. An edge pointing away (towards) a node represents a downstream (upstream) contract for the agent represented by this node. This restriction is natural in many markets – for example, a power station cannot sell electricity directly to a household.

Each firm i has a strict preference relation P^i over sets of contracts Ω_i , which involve i . For any $\Omega' \subseteq \Omega$ define the chosen set for firm i as

$$C^i(\Omega') \equiv \max_{P^i} \{\Omega'' \subseteq \Omega' | \omega \in \Omega'' \Rightarrow i \in \omega_I\} \quad (3.8)$$

In order to define preferences concisely, let us introduce a little more notation. For any

Figure 3.3: Prohibited contract structure



Contractual network: nodes represent agents; edges represent contracts.

$B'_i \subseteq B_i$ and $S'_i \subseteq S_i$, define the chosen set

$$C_b^i(B'_i|S'_i) \equiv \{\omega \in C^i(\{\omega' \in B'_i\} \cup \{\omega' \in S'_i\}) | \omega_b = i\} \quad (3.9)$$

which is the set of downstream contracts i chooses as a buyer when i has access to downstream contracts B'_i and upstream contracts S'_i . Analogously, define

$$C_s^i(S'_i|B'_i) \equiv \{\omega \in C^i(\{\omega' \in B'_i\} \cup \{\omega' \in S'_i\}) | \omega_s = i\} \quad (3.10)$$

Hence, we can define rejected sets of contracts $R_b^i(B'_i|S'_i) \equiv B'_i \setminus C_b^i(B'_i|S'_i)$ and $R_s^i(S'_i|B'_i) \equiv S'_i \setminus C_s^i(S'_i|B'_i)$. An *allocation* $\Psi \subseteq \Omega$ is a set of contracts.

Now, we can come to the assumption on preferences, which, in the setting of finite contracts, is an analogue of gross substitutes and complements.

Definition 3.6. Preferences of $i \in I$ over contracts in Ω are *fully substitutable* if for all $B''_i \subseteq B'_i \subseteq \Omega$ and $S''_i \subseteq S'_i \subseteq \Omega$ they are:

1. *Same-side substitutable*:

(a) $R_b^i(B''_i|S'_i) \subseteq R_b^i(B'_i|S'_i)$

(b) $R_s^i(S''_i|B'_i) \subseteq R_s^i(S'_i|B'_i)$

2. *Cross-side complementary*:

$$(a) R_b^i(B'_i|S'_i) \subseteq R_b^i(B'_i|S''_i)$$

$$(b) R_s^i(S'_i|B'_i) \subseteq R_s^i(S'_i|B''_i)$$

Contracts are fully substitutable if a firm regards any two upstream or any two downstream contracts as substitutes, but an upstream and a downstream contract as complements.

3.5.2 Cycles, chains, and stability

As we mentioned above, our model does not restrict the contractual network to a supply chain. Instead we allow firms to form contractual cycles.

Definition 3.7. A set of contracts $\{\omega^1, \dots, \omega^M\} \subseteq \Omega$ is a *cycle* if $\omega_b^1 = \omega_s^2, \omega_b^2 = \omega_s^3, \dots, \omega_b^{M-1} = \omega_s^M$ and $\omega_b^M = \omega_s^1$.

A cycle is odd if $M > 2$ is odd. An illustration of an even contractual cycle is, in fact, provided in Figure 3.4. In this example of a fairly complex market structure, there are four segments in \mathcal{C} (which are labelled and marked in different colours) and one even contract cycle marked in grey.


To formalise our main solution concept for contractual networks – chain stability – we first define a chain, which is a set of contracts that contains no cycles or sub-cycles.

Definition 3.8. A set of contracts $\{\omega^1, \dots, \omega^M\} \subseteq \Omega$ is a *chain* if

1. $\omega_b^m = \omega_s^{m+1}$ for all $m = 1, \dots, M - 1$.
2. $\omega_s^m = \omega_s^n$ implies $n = m$.
3. $\omega_b^M \neq \omega_s^1$

Finally, we define chain stability (Ostrovsky, 2008), which is an analogue of pairwise stability (Roth, 1984b) in a classical matching market.

Figure 3.4: Even cycles in the contractual network



Contractual network: nodes represent agents; edges represent contracts.

Definition 3.9. An allocation $\Psi \subseteq \Omega$ is *chain-stable* if it is:

1. *Individually rational*: for all $i \in I$, $C^i(\Psi) = \Psi_i$.
2. *Chain unblocked*: There is no chain $Z \subseteq \Omega$, such that $Z \not\subseteq \Psi$ and for all $i \in Z_I$, $Z_i \subseteq C^i(\Psi \cup Z)$.

3.5.3 Result

We can now present the second main result of the chapter, which has a lot in common with the first. In the proof, we show that any set of contracts in a segmented market without odd cycles (e.g. Figure 3.4) can be represented by an appropriate undirected bipartite graph where nodes represent contracts and edges represent firms (e.g. Figure 3.2).

Theorem 3.2. *Suppose that in a segmented market the set of contracts Ω has no odd cycles and that preferences of I are fully substitutable. Then a chain-stable contract allocation Ψ exists.*

Proof. If the set of contracts Ω does not contain any cycles, then the result follows immediately as a corollary of Theorem 1 in Ostrovsky (2008). Suppose, it contains no


odd contractual cycles. Consider a directed graph G , whose arcs (directed edges) are the contracts $\{\omega^1, \dots, \omega^N\} \subseteq \Omega$ and whose nodes are the agents who are party to each contract. Since Ω does not form an odd cycle, there is a bipartite graph G' whose nodes (that represent contracts i.e. arcs of Ω) can be divided into two disjoint sets Ω^B and Ω^S such that every edge (that represents a node from Ω i.e. an agent) connects node in Ω^B to the node(s) in Ω^S only if an agent is involved in both contracts. Consider any agent $i \in I$. Without loss of generality, assume that B_i is a set of nodes in Ω^B (call it Ω_i^B) and S_i is a set of nodes in Ω^S (call it Ω_i^S , for symmetry). Agent i has fully substitutable preferences over some $B_i \subseteq \Omega^B$ and some $S_i \subseteq \Omega^S$. Now consider three new types of agents: buyers, sellers, and intermediaries. Buyers and sellers only sign contracts with the intermediaries. Buyers are potential parties to all contracts in Ω^B only and sellers are potential parties to all contracts in Ω^S only and the intermediaries are parties to $\Omega^B \cup \Omega^S$. The preferences over $T \cap (\Omega_i^B \cup \Omega_i^S)$ for any $T \subseteq \Omega^B \cup \Omega^S$ for all three types of agents must match those of $i \in I$. Hence, the buyers', the sellers' and the intermediaries' preferences are fully substitutable over Ω^B and Ω^S . Since the buyers and the sellers only sign contracts with the intermediaries, every contract allocation contains no cycles and this satisfies a supply-chain structure. Once again, the result follows from Theorem 1 in Ostrovsky (2008). \square

Hatfield and Kominers (2012) showed that whenever preferences are fully substitutable and contracts are acyclic chain stability is equivalent to *stability* – i.e. immunity of deviations by sets of firms.

Definition 3.10. An allocation $\Psi \subseteq \Omega$ is *stable* if it is:

1. *Individually rational*: for all $i \in I$, $C^i(\Psi) = \Psi_i$.
2. *Unblocked*: There is no non-empty set of contracts $Z \subseteq \Omega$, such that $Z \not\subseteq \Psi$ and for all $i \in Z_I$, $Z_i \subseteq C^i(\Psi \cup Z)$.

Figure 3.5: Example of a network that is chain-stable, but not stable



Contractual network: nodes represent agents; edges represent contracts.

In general chain stability is a weaker stability notion than stability, and chain-stable contract allocations are not necessarily stable. The chain-stable contract allocation may, of course, be an empty set of contracts. The following example from Hatfield and Kominers (2012) illustrates both of these points. Consider three contracts ω^1 , ω^2 and ω^3 . Assume that $i = \omega_b^1$, $j = \omega_s^1 = \omega_s^2 = \omega_b^3$, and $k = \omega_b^2 = \omega_s^3$ (see Figure 3.5). Preferences are fully substitutable in the following way: $P^i : \{\omega^1\} \succ \emptyset$, $P^j : \{\omega^1, \omega^3\} \succ \{\omega^2, \omega^3\} \succ \emptyset$, $P^k : \{\omega^2, \omega^3\} \succ \emptyset$, but only form an even cycle. Hence, a chain-stable allocation exists: $\Psi = \emptyset$.⁹ There is, however, no stable allocation. Our results do not contradict Theorem 7 in Hatfield and Kominers (2012) since Theorem 3.2 only considers the existence of chain-stable outcomes.

3.6 Discussion

We now discuss how our framework can be applied to a variety of existing models.

3.6.1 Trading networks

Hatfield et al. (2013) consider a model of trading networks in which agents' preferences are quasilinear and fully substitutable over the contracts they are involved in.

⁹Chain-stable outcomes in the absence of odd cycles also correspond to an appropriately defined fixed-point operator (Ostrovsky, 2008).

They show that competitive equilibrium (with and without personalised prices) exists in their model in the presence of arbitrary contractual cycles. However, the results of Hatfield et al. (2013) do not contradict our Theorem 3.1 since they do not actually permit SC structures with odd cycles illustrated in section 3.3.4.1. To see this, note that because of quasilinearity all trades can be divided into disjoint subsets of upstream and downstream trades for every agent. Now relabel every upstream and downstream part of each trade as a good. No agent can regard two downstream (or upstream) trades as complements, so any cycle in the SC structure must be even.¹⁰

Drexler (2013) generalised this model to the case of “full substitutes and complements” by allowing agents to view both their upstream and downstream contracts as gross substitutes and complements. Our results for the trading economy show that this can be made even more general – we can allow a SC structure with no odd cycles on upstream and downstream contracts. Moreover, the properties of competitive equilibria in trading networks, described Hatfield et al. (2013), continue to hold whenever preferences have a SC structure with no odd cycles. In particular, the allocation supported by a competitive equilibrium is *strong group stable* (a strengthening of *stability*, see their paper and Chapter 2 of this thesis for the definition) and in the core.

¹⁰In a footnote on p. 18, Hatfield et al. (2013) say

Note that our network setting is clearly more general, as it cannot be embedded in the setting of Sun and Yang (2006). E.g., consider a simple market with three agents i, j , and k , where i can sell trades to both j and k , j can sell trades only to k , k cannot sell trades to anyone, and all agents’ preferences are fully substitutable. In this market, it is impossible to separate trades into two groups S_1 and S_2 in such a way that every agent views trades in one group as substitutes and views trades in different groups as complements.

This reasoning is somewhat misleading. Rather than having preferences over trades, agents have preferences over roles (buyer or seller) that they play in these trades. In fact, since preferences are quasilinear, we can separate the possible downstream trades for each agent into S_1 and the possible upstream trades for each agent into S_2 and recover the Sun and Yang’s framework.

3.6.2 Importance of quasilinearity

The proofs illustrate both the differences and similarities between GSC and fully substitutable preferences. Conceptually, fully substitutable preferences for a buyer, a seller and an intermediary are an analogue of GSC in the case of contracts without continuous transfers. However, subtle, yet important, differences arise when one considers different notions of stability of contract allocations. Hatfield et al. (2013, p. 18) point out that:

for contractual sets that allow for continuous transfers, in the presence of quasilinearity, supply chain structure is not necessary for the existence of stable outcomes, although full substitutability is. It is an open question why the presence of a continuous numeraire can replace the assumption of a supply chain structure in ensuring the existence of stable outcomes.

In this chapter, we have shown that supply chain structure is not necessary for the existence of chain-stable outcomes even without quasilinearity. Competitive equilibrium exists in arbitrary trading networks because they preclude odd cycles in the SC structure of preferences.

3.6.3 Auction

Sun and Yang (2008, 2009) introduced a tâtonnement process – called a double-track auction – that finds the competitive equilibrium allocation and prices when agents' preferences satisfy GSC. The process is powerfully simple. The seller starts off with low prices for goods in S_1 and with high prices in S_2 . At every stage of the process, the seller asks agents to report their demands. The prices of over-demanded goods are increased and prices of under-demanded goods are decreased. Without stating all the details, the example in section 3.3.4.1 and the proof of Theorem 3.1 make it clear how the double-track auction can be generalised to find competitive equilibrium allocation prices for any trading economy in which agents' preferences have a SC structure without odd cycles.

The seller starts off with low prices for goods in (nodes) S^A and high prices for goods in (nodes) S^B and follows the price adjustment rules of the double-track auction as if the goods were contained in S_1 and S_2 . Whenever there are odd cycles in the SC structure of preferences, the seller would have to start off the auction with high (low) prices for goods regarded as complements by some agents and the descending (ascending) auction may not find the market clearing prices (Sun and Yang, 2008, Section 3.2). Similarly, we can easily adapt the strategy-proof and efficient dynamic mechanism based on the double-track auction for preferences that satisfy SC structure with no odd cycles (Sun and Yang, 2008, Section 4).

3.6.4 Algorithm for testing full substitutability

As Figure 3.2 shows, even complex SC structures may have no odd cycles. There exist linear time algorithms which can check whether a graph representing a SC structure is bipartite (i.e. contains no odd cycles). Therefore, it is easy to test the presence of odd cycles provided the SC structure of the market is well known.

However, it is not clear whether there exists a polynomial time algorithm to test full substitutability. Hatfield et al. (2012) showed that this algorithm (polynomial in the number of contracts and the preference list length) exists for simple substitutable preferences (Hatfield and Milgrom, 2005) and further research may test whether this can be carried over to fully substitutable preferences in segmented markets without odd contractual cycles.

3.7 Conclusion

We showed that competitive equilibrium exists in a trading economy when agents' preferences have a substitutes and complements structure with no odd cycles. We then applied this intuition to contractual networks. We proved that a chain-stable contract allocation may exist even when contractual cycles are present. The results

of this chapter can, of course, be immediately used in market design applications, such as auctions or matching markets. We illustrated the obvious similarities and subtle differences in preferences, which satisfy GSC and full substitutability. It would be interesting to see whether by strengthening the assumptions on preferences over contracts chain-stable outcomes can exist whenever any contractual cycle is present. Additionally, it is important to examine what sort of contracts can be represented by quasilinear preferences. This would allow us to understand the relationship between quasilinearity, competitive equilibrium, and different notions of stability more deeply. We leave both of these topics for further research.

3.8 Appendix

Absence of odd cycles is not necessary for the existence of a competitive equilibrium. We claimed that absence of odd cycles in the SC structure is only a sufficient condition for stability. As our example above showed, there are valuations for agents with the same preferences such that no competitive equilibrium in the economy exists. The example illustrated in Table 3.3 below shows that lack of odd cycles in SC structure is not necessary for equilibrium existence.

Table 3.3: Competitive equilibrium with an odd cycle in the SC structure

Ψ	$u_i(\Psi_i)$	$u_j(\Psi_j)$	$u_k(\Psi_k)$
\emptyset	0	0	0
$\{\omega_1\}$	2	0	4
$\{\omega_2\}$	2	1	0
$\{\omega_3\}$	0	1	4
$\{\omega_1, \omega_2\}$	5	0	0
$\{\omega_2, \omega_3\}$	0	3	0
$\{\omega_3, \omega_1\}$	0	0	9
$\{\omega_1, \omega_2, \omega_3\}$	5	3	9

Price vector $p^* = (4.5, 2, 4.5)$ supports a competitive equilibrium allocation $\Pi^* = \{\{\Psi_i, \Psi_j, \Psi_k\} = \{\{\omega_2\}, \emptyset, \{\omega_1, \omega_3\}\}\}$. We find this by inspection since we cannot apply the double-track auction. The following proposition is trivial to prove and follows from our examples in section 3.3.4.

Proposition 3.1. *For any trading economy $\mathcal{E} = \{\Omega, (u_i, i \in I)\}$, if agents' preferences have a SC structure on \mathcal{S} and form odd SC cycles, there exist agents' valuations over bundles with the same SC structure on \mathcal{S} , such that no competitive equilibrium exists.*

References

- Baldwin, E. and P. Klemperer (2013). Tropical geometry to analyse demand. Discussion paper, University of Oxford, Department of Economics.
- Bikhchandani, S. and J. W. Mamer (1997). Competitive equilibrium in an exchange economy with indivisibilities. *Journal of Economic Theory* 74(2), 385–413.
- Chung, K.-S. (2000). On the existence of stable roommate matchings. *Games and Economic Behavior* 33(2), 206–230.
- Crawford, V. P. and E. M. Knoer (1981). Job matching with heterogeneous firms and workers. *Econometrica* 49(2), 437–450.
- Drexl, M. (2013). Substitutes and complements in trading networks. Technical report, Mimeo.
- Gudmundsson, J. (2013, May). Cycles and third-party payments in the partnership formation problem. Working paper 2013:16, Lund University, Department of Economics, School of Economics and Management.
- Gul, F. and E. Stacchetti (1999). Walrasian equilibrium with gross substitutes. *Journal of Economic Theory* 87(1), 95–124.
- Hatfield, J. W., N. Immorlica, and S. D. Kominers (2012). Testing substitutability. *Games and Economic Behavior* 75(2), 639–645.
- Hatfield, J. W. and S. D. Kominers (2012). Matching in networks with bilateral contracts. *American Economic Journal: Microeconomics* 4(1), 176–208.
- Hatfield, J. W., S. D. Kominers, A. Nichifor, M. Ostrovsky, and A. Westkamp (2013). Stability and competitive equilibrium in trading networks. *Journal of Political Economy*, forthcoming.
- Hatfield, J. W. and P. Milgrom (2005). Matching with contracts. *American Economic Review* 95(4), 913–935.
- Kelso, A. S. and V. P. Crawford (1982). Job matching, coalition formation, and gross substitutes. *Econometrica* 50(6), 1483–1504.
- Milgrom, P. (2000). Putting auction theory to work: The simultaneous ascending auction. *Journal of Political Economy* 108(2), 245–272.
- Milgrom, P. and J. Roberts (1990). The economics of modern manufacturing: Technology, strategy, and organization. *American Economic Review* 80(3), 511–528.

- Milgrom, P. and B. Strulovici (2009). Substitute goods, auctions, and equilibrium. *Journal of Economic Theory* 144(1), 212–247.
- Ostrovsky, M. (2008). Stability in supply chain networks. *American Economic Review* 98(3), 897–923.
- Roth, A. E. (1984b). Stability and polarization of interests in job matching. *Econometrica* 52(1), 47–58.
- Shioura, A. and Z. Yang (2013). Equilibrium, auction, multiple substitutes and complements. Discussion Paper 13/17, University of York.
- Sun, N. and Z. Yang (2006). Equilibria and indivisibilities: Gross substitutes and complements. *Econometrica* 74(5), 1385–1402.
- Sun, N. and Z. Yang (2008). A double-track auction for substitutes and complements. Discussion Paper 656, Kyoto Institute for Economic Research.
- Sun, N. and Z. Yang (2009). A double-track adjustment process for discrete markets with substitutes and complements. *Econometrica* 77(3), 933–952.
- Sun, N. and Z. Yang (2011). A general competitive trading network with indivisibilities and gross substitutes and complements. Mimeo.
- Talman, D. and Z. Yang (2011). A model of partnership formation. *Journal of Mathematical Economics* 47(2), 206–212.
- Tan, J. (1991). A necessary and sufficient condition for the existence of a complete stable matching. *Journal of Algorithms* 12(1), 154–178.
- Westkamp, A. (2010). Market structure and matching with contracts. *Journal of Economic Theory* 145(5), 1724–1738.

Chapter 4

Auctions for complements: an experimental investigation

Abstract

We experimentally examine the effects of exposure and package bidding in sealed-bid auctions. We focus on four auctions: first-price, Vickrey, as well as Vickrey-Nearest Rule and Reference Rule. In our setting, three bidders (two local and one global) bid on two items. We find that the first-price auction is revenue-dominant without any loss of efficiency across all exposure and package treatments. Exposure harms the exposed global bidders' profits in the first-price, but not in other auctions. We do not find any substantial differences between the Vickrey-Nearest Rule and the Reference Rule. Local bidders bid truthfully in the Vickrey auction. However, there is no evidence of threshold effects. Finally, we find that in package settings, bidders fail to take into account the effect of their own bids on their payments.

Keywords: auctions, complementarities, exposure, package bidding, experiments.

JEL Classification: *D44, C91, C92*

4.1 Introduction

What is the best way to sell indivisible goods? An auction is a popular and convenient market design instrument for the task.¹ In some cases, economic theory provides clear guidance on which auction to use. When there is one item on sale, a first- or second-price sealed-bid auction with an optimal reserve price maximises expected revenue for the seller (Myerson, 1981) and the second-price auction is efficient. In the case of substitutable goods, Kelso and Crawford (1982) showed that bidding good by good in a simple dynamic auction ensures convergence to an efficient allocation; in fact, the celebrated sealed-bid Vickrey (1961) auction is also efficient and delivers seller revenues and bidder payments which are in the core (Milgrom, 2004, Theorem 8.5).²

In many auction settings, some items on sale are complements. This means that for some bidders the joint valuation of the items is higher than sum of the valuations of the individual items. Complementarities may be asymmetric across bidders and arise in many real-world auctions – for mobile spectrum, for airport takeoff and landing slots, and for bus routes (Cramton et al., 2006) – due to technological constraints, increasing returns to scale, and so on. However, Kelso and Crawford (1982) showed that, whenever complementarities are present, the core and the set of competitive equilibria might be empty, so bidding good by good no longer guarantees to produce an efficient outcome. As we argue below, complementarities create further challenges for sealed-bid auction designers. In this chapter, we experimentally examine what sealed-bid auctions perform well in various settings when complementarities are present.

If items are complements but bidders can only bid on individual items, complementarities may leave these bidders *exposed*. For example, suppose there is an auction for

¹Milgrom (2004), Klemperer (2004), and Krishna (2009) provide excellent surveys of recent auction literature and practice.

²These results hold provided there is one unit of each good. The core is the set of coalition-proof outcomes.

a pair of trousers and a jacket.³ Bidder 1 only wants to buy the trousers and Bidder 2 only wants to buy the jacket. In this chapter, we will refer to Bidders 1 and 2 as *local* bidders. Bidder 3, however, wants to buy both items and values either item at zero (perhaps because he is going to a job interview). We call Bidder 3 the *global* bidder. Therefore, if Bidder 3 wins either item without the other, he will be exposed to the risk of negative profit. In order to avoid this situation, the bidder might bid cautiously. Such cautious bidding is likely to lead to an inefficient allocation of items on sale and low revenues for the auctioneer.

There are various ways to deal with the exposure problem. One solution is to *bundle* the items. The seller could bundle the trousers and the jacket into a single item (call it a suit) and offer it to the bidders. The seller needs to have considerable information about bidders' valuations of items in order to bundle them correctly. In the case where all bidders are global and want to buy the suit, this is indeed *optimal* (i.e. revenue-maximising) for the seller (Gale, 1990). Yet bundling can create inefficiencies if the bidders valuations are asymmetric e.g. in the case that either Bidder 1 or Bidder 2 wins the suit, but their joint valuation of the trousers and the jacket is higher than that of Bidder 3. The Vickrey auction will nevertheless be efficient subject to the bundling constraint. However, with more than two global bidders, bundling may not be optimal in the Vickrey auction (Krishna, 2009, p. 232).

Another solution to the exposure problem is to allow bidders to bid on packages as well as on individual items. If Bidder 3 could bid on the package of items containing the trousers and jacket, and Bidders 1 and 2 could bid on the individual items, we could, in principle, recover full efficiency. Such *package*, or *combinatorial*, auctions are increasingly popular in industry and government procurement. They found wide-

³Other analogous examples are: spectrum licences to operate in the West and the East regions of a country; 10am take-off slot at London Heathrow and 11am landing slot at Edinburgh Airport; contracts to run day and night bus services on the same route.

ranging applications from freight transportation services (Sheffi, 2004) to determining providers of school meals in Chile (Epstein et al., 2004), and cleaning services in Sweden (Lunander and Lundberg, 2012).

The efficient package auction in our example is still the Vickrey auction. In the package Vickrey auction, the bidders' payments equal to the opportunity cost of their presence in the auction. Since their payments, conditional on winning, do not depend on their bids, it is a (weakly) dominant strategy for every bidder to report their true valuation of every package. However, in the presence of complementarities, the Vickrey auction suffers from significant drawbacks.⁴ The most important one is its proneness to collusion resulting in very low seller revenues. Suppose that Bidders 1 and 2 both submit bids on individual items, which are equal to (or higher than) the package bid of Bidder 3. In this case, Bidders 1 and 2 win their items (because the sum of their bids is strictly higher than the bid of Bidder 3) and their Vickrey payments are zero (because the opportunity cost of their presence in the auction is zero). Hence, seller revenues are also zero.⁵ The outcome of the Vickrey auction in this case is not in the core i.e. it can be blocked by the seller selling the suit to Bidder 3 at a small price ϵ making both of them better off.

In order to avoid collusive behaviour, the seller could run a first-price combinatorial auction in which winners pay their bids. However, in the presence of complementarities, first-price combinatorial auctions may suffer from the well-known *threshold* problem (Bykowsky et al., 2000, Milgrom, 2000). Consider, for example, a first-price combinatorial auction for the jacket and the trousers. Both Bidder 1 and Bidder 2 know that if the sum of their bids exceeds that of Bidder 3, they will win their respective items.

⁴Milgrom (2004, pp. 56-61) provides for an exhaustive list of the drawbacks of the Vickrey auction in the presence of complements.

⁵Indeed, zero revenues may arise without any collusion and even when bidders bid truthfully. If both local bidders' valuations of the individual items are simply higher than the global bidder's valuation of the of package, then truthful bidding also produces zero revenues in the Vickrey auction.

However, they both face incentives to *free-ride* on the higher bid of the other bidder since their payments, whenever they win, equal their bids. This miscoordination may result in lower local bidders' bids and, hence, inefficiency and lower seller revenues.

Intending to find a compromise between the possibility of low revenues in the Vickrey auction and the inefficiency of the first-price auction, a new class of minimum-revenue core-selecting (MRCS) auctions has been proposed in a number of papers in recent years. In this chapter, we consider the two most prominent MRCS auction designs, namely the *Reference Rule* (Erdil and Klemperer, 2010) and the *Vickrey-Nearest Rule* (Day and Cramton, 2008, 2012). These auction rules deal with low revenues – by demanding that winning bidder payments are in the bidder-optimal core – as well minimising the incentives to deviate from truthful bidding that may harm efficiency.⁶

In an important theoretical contribution, Ausubel and Baranov (2010) provided – whenever it was possible – analytical and numerical solutions for equilibrium bidding strategies, revenue, and efficiency rankings between the first-price, Vickrey, Reference Rule and the Vickrey-Nearest Rule auctions in a setting with two ex ante symmetric local bidders and one global bidder. In their setting, the local bidders can only bid on the items that they are interested in and the global bidder can only submit a bid on the package of two items. This chapter builds on the contribution by Marszalec (2011), who ran a series of within-subject auction experiments, which tested Ausubel and Baranov's theoretical results. Although Ausubel and Baranov (2010) showed that, in equilibrium, the Vickrey auction was fully efficient *and* revenue-dominated all the other auctions, Marszalec (2011) found that:

1. Vickrey auction turned out to be the least efficient out of the four auctions as local bidders attempted to collude despite anonymous random matching. However, the success rate for collusion was low.

⁶Individual incentives to deviate from truthful bidding are, of course, entirely absent in the Vickrey auction. As we showed above, group incentives to deviate may still be present.

2. When bidding restrictions were introduced to prevent collusive behaviour, all auctions were equally efficient.
3. First-price auction dominated all the other auctions in revenue with and without bidding restrictions rejecting the theoretical results of Ausubel and Baranov (2010).⁷
4. The hypotheses of equilibrium and “sophisticated” bidding were rejected in all auctions.⁸

In this chapter, we run 16 treatments i.e. four treatments for each of the four auctions: first-price, Vickrey, Vickrey-Nearest Rule, and Reference Rule. We also use a two-item, three-bidder (two local and one global) setting. In the control (Bundle) treatment, global bidder can only bid on the package following Marszalec (2011). In two treatments (Exposed and Components), we leave the global bidder exposed by restricting his bids to individual items. These two treatments have different levels of exposure. In another (Combinatorial) treatment, we allow the global bidder to bid both on the package and on the individual items. We can summarise our findings as follows:

1. First-price auction is revenue-dominant across all treatments without much loss of efficiency. There is no difference in revenue between the other three auctions.
2. Exposure affects profits of the exposed global bidders only in the first-price auction.

⁷In a separate project with Marszalec, we also found that, whenever bidding restrictions are in place, results 2 and 3 above hold are robust for competition (up to six bidders) and bidder information treatments.

⁸The behavioural notion of “sophisticated bidding” suggests that players are, in fact, best-responding to the actual play of their opponents, rather than to Bayes-Nash equilibrium strategies, since the opponents may not be following equilibrium play. If this were the case, then the best response may indeed not be the same as suggested by theoretical equilibrium strategies. However, this hypothesis is firmly rejected by the experimental data; subjects could have gained by unilaterally deviating to Bayes-Nash equilibrium strategies.

3. Global bidders do not take into account the effect of their own bids on their payments in the Combinatorial setting.
4. Local bidders bid truthfully in the Vickrey auction.
5. There are no threshold effects in any auction.
6. We find no pronounced differences between the two MRCS auctions.

We proceed as follows. Section 4.2 discusses the relevant theoretical and experimental literature. In section 4.3, we discuss the experimental design, provide the theoretical background, and present the hypotheses we will test. Section 4.4 describes the experimental procedures. Section 4.5 presents and discusses the results, while section 4.6 concludes. The Appendix contains experiment instructions, further details about the experiment as well as tables and figures omitted from the main text.

4.2 Related literature

4.2.1 Theory

Theoretical interest in auctions where bidders view some items as complements arose during the U.S. Personal Communications Service (PCS) spectrum rights auctions in the early 90s (McMillan, 1994). The Federal Communications Commission (FCC) noted that, due to economies of scale, complementarities for licenses could be so high that it might even be efficient to allocate licenses in all trading areas to one bidder. However, the FCC decided against running a combinatorial auction for computational reasons.⁹ To this end, Krishna and Rosenthal (1996) examine a second-price, sealed-bid auction for two items in which there are local and global bidders (cf. our Exposed and Components treatments in the Vickrey/Vickrey-Nearest Rule/Reference Rule auctions). They

⁹In general, the integer-programming problem that determines the winners in a combinatorial auction is NP-hard i.e. a polynomial time algorithm to solve it is not known to exist.

find that global bidders bid more aggressively when the synergies are higher, but less aggressively when they face more competition from other global bidders. Levin (1997) considered optimal mechanisms to sell two items to multiple global bidders. He showed that most standard auctions will not be optimal and bundling is optimal only when bidders are perfectly symmetric.

In recent years, there have been several attempts to address the deficiencies of standard combinatorial auction formats in the presence of complementarities. The most influential class of auctions that emerged in theory and practice are minimum-revenue (bidder-optimal) core-selecting (MRCS) auctions. Day and Raghavan (2007) and Day and Milgrom (2008) attempted the first systematic study of MRCS auctions in a complete information environment. The motivation for selecting payments in the minimum-revenue core (MRC) is powerful: among all the core outcomes, they minimised the total sum of unilateral incentives to deviate from truth-telling. Moving to incomplete information environments considered in this chapter, Beck and Ott (2011) studied a setting in which local bidders could also place package bids and found that in equilibrium these bids may be above true values. Goeree and Lien (2012) showed that any MRCS auction the outcome of which does not coincide with the Vickrey auction is not Bayesian incentive-compatible. Therefore, in general incomplete information settings, incentives to deviate from truth-telling will always remain in package MRCS auctions with complements.¹⁰

Since the MRC is not usually a unique point, Day and Cramton (2008, 2012) proposed selecting the payments on the MRC that minimise the Euclidian distance from the Vickrey payments. In practice, many MRCS auctions have been run that way, most recently in the sale of 4G spectrum licences in the United Kingdom. However, Erdil and Klemperer (2010) pointed that it is possible to reduce the sum of the bidders'

¹⁰See related results by Lamy (2010), Sano (2011) and Hafalir and Yektaş (2012).

marginal incentives to deviate from truth-telling by selecting payments that are closest to a pre-determined reference point on the MRC. In order to do this, the seller could assign the reference point by fixing the relative prices for goods paid by the winners. These relative prices could be determined using existing market information or from equity considerations. We discuss both of these auctions in more detail in section 4.3.

4.2.2 Experiments

Roth (2008) elegantly summarises the role of experiments in auctions and other market design applications:

Experiments have multiple uses in market design, not only for investigation of basic phenomena, and smallscale testing of new designs, but also in the considerable amount of explanation, communication and persuasion that must take place before designs can be adopted in practice.

Indeed, Ledyard et al. (1997) and Plott (1997) discuss the role that experiments played in designing and fine-tuning the PCS auction. Earlier, Dyer et al. (1989) provided evidence that bidding behaviour of university students and experts (business executives who participate in construction auctions) is very similar in an experimental lab.

Kagel and Levin (2011) provide an exhaustive review of the literature on experiments in auctions. Most of experiments that involved selling complementary items tested non-package dynamic auctions, such as the SMR auction. Several papers compared the SMR auction to various alternatives in settings with large number of items (6-10) and bidders (3-5) (Kwasnica et al., 2005, Brunner et al., 2010, Kagel et al., 2010, 2012). Goeree and Holt (2010) also compared several simple, dynamic, combinatorial auctions in which bidders were interested in substitutable and complementary goods. We will show that these sort of large-scale settings will be too demanding for experimental subjects facing either of the sealed-bid MRCS auctions considered in this chapter or even simple Vickrey auctions.

There have been fewer experiments testing sealed-bid auctions with complements. Kagel and Levin (2005) tested a non-package, second-price, sealed-bid auction similar to the one introduced by Krishna and Rosenthal (1996) against an ascending clock auction. They found, unsurprisingly, that subjects bid closer to optimal bidding strategies in the dynamic auction. Chen and Takeuchi (2010) found that a version of an ascending-price generalised Vickrey (iBEA) auction outperforms the simple Vickrey auction in a setting with three bidders and four complementary items.

However, since sealed-bid combinatorial auctions are often used in the final allocation phase of spectrum auctions, understanding their properties experimentally is of both academic and policy interest. To this end, Chernomaz and Levin (2012) have done a series of experiments testing the effect of exposure and threshold bidding in first-price auctions with complements. Their two-item, three-bidder setting is similar to ours, but their local bidders assign identical value to the items. They find that when bidders are exposed package bidding can improve efficiency. However, by introducing the threshold problem, package bidding always lowers revenues. Most of our findings for the first-price auctions are qualitatively consistent with Chernomaz and Levin's, however, we do not find that package bidding reduces revenues. Moreover, the effects of package bidding are different in the MRCS and Vickrey auctions.

Finally, apart from the contribution by Marszalec (2011), there have been no experimental results that tested various MRCS and classic auctions head to head. This chapter also fills this gap.

4.3 Theoretical background, experimental design, and hypotheses

4.3.1 Model set-up

We consider a simple setting, in which a seller wants to sell two items $K = \{x, y\}$ to three bidders $I = \{1, 2, 3\}$. There are four possible packages in total: $\{\emptyset, \{x\}, \{y\}, \{x, y\}\}$. Each bidder $i \in I$ assigns a value to each package $\mathbf{v}_i = (v_i(S))_{S \subseteq K}$. We assume that $v(\emptyset) = 0$. Two *local* bidders $\{1, 2\}$ only want to buy individual items. Bidder 1 wants to buy item x and Bidder 2 wants to buy item y , so we denote their valuations of these items by v_1^x and v_2^y , respectively.¹¹ v_1^x and v_2^y are independently drawn from a uniform distribution with support $[0, 100]$. We assume that Bidder 1 has no use for item y and Bidder 2 has no use for item x , therefore $v_1^y = v_2^x = 0$, $v_1^{xy} = v_1^x$, and $v_2^{xy} = v_2^y$. Global Bidder 3 is interested in buying both items x and y . His valuations of acquiring both items, v_3^{xy} , and of the individual components, v_3^x and v_3^y , depend on the treatment.

Agents' valuations are private knowledge. Utility is quasilinear so bidder profit is $u_i(S) = v_i(S) - p_i(S)$ for any $S \subseteq \{x, y\}$ that a bidder $i \in I$ wins. Finally, we assume that the seller values the items at 0.

4.3.2 Treatments

Every auction is run under four conditions, which we refer to as *treatments*. These four treatments vary global Bidder 3's valuation function *and* affect whether he is allowed to place bids on the package and/or on the individual items.

- *Combinatorial*
 - Bidder 3 can place bids on the individual items x and y *and* on the package xy .

¹¹For the brevity of notation, we will always write $v_1^x \equiv v_1(\{x\})$, $v_1^{xy} \equiv v_1(\{x, y\})$ and so on.

- Bidder 3’s valuations for the individual items are drawn from a uniform distribution on $[0, 100]$. Bidder 3’s valuation of acquiring both items x and y is the sum of the valuations for the individual items multiplied by 1.5, but at most 200. Therefore, $v_3^x, v_3^y \sim U(0, 100)$ and $v_3^{xy} = \min\{1.5 \times (v_3^x + v_3^y), 200\}$.
- Bidder 3 can win both items (and obtain valuation v_3^{xy}) either by winning the items with individual bids on the items or by placing a winning bid on the package xy .

- *Components*

- Bidder 3 can *only* place bids on the individual items x and y .
- Bidder 3’s valuations for the individual items are drawn from a uniform distribution on $[0, 100]$. Bidder 3’s valuation of acquiring both items x and y is the sum of the valuations for the individual items multiplied by 1.5, but at most 200. Therefore, $v_3^x, v_3^y \sim U(0, 100)$ and $v_3^{xy} = \min\{1.5 \times (v_3^x + v_3^y), 200\}$.
- Bidder 3 can only win both items (and obtain valuation v_3^{xy}) by winning the items with individual bids on the items.

- *Exposed*

- Bidder 3 can *only* place bids on the individual items x and y .
- Bidder 3’s valuations for the individual items are zero. Bidder 3’s valuation of acquiring both items x and y is the sum of two numbers drawn independently from a uniform distribution on $[0, 100]$, multiplied by 1.5, but at most 200. Therefore, $v_3^x = v_3^y = 0$ and $v_3^{xy} = \min\{1.5 \times (U(0, 100) + U(0, 100)), 200\}$.
- Bidder 3 can only win both items (and obtain valuation v_3^{xy}) by winning the items with individual bids on the items.

- *Bundle*

- Bidder 3 can *only* place bids on the package xy .
- Bidder 3’s valuations for the individual items are zero. Bidder 3’s valuation of acquiring both items x and y is the sum of two numbers drawn independently from a uniform distribution on $[0, 100]$, multiplied by 1.5, but at most 200. Therefore, $v_3^x = v_3^y = 0$ and $v_3^{xy} = \min\{1.5 \times (U(0, 100) + U(0, 100)), 200\}$.
- Bidder 3 can only win both items (and obtain valuation v_3^{xy}) by placing a winning bid on the package xy .

The Bundle treatment is effectively a control treatment. As in the theoretical settings of Ausubel and Baranov (2010), Erdil and Klemperer (2010) and Day and Cramton (2012) as well as in the experiments by Marszalec (2011), Bidder 3 can only bid on the package xy . Although Bidder 3’s valuation of the package is drawn from a different distribution than in the latter paper, we also force the support to be on $[0, 200]$. The complementarity coefficient of 1.5 is also used by Chernomaz and Levin (2012).

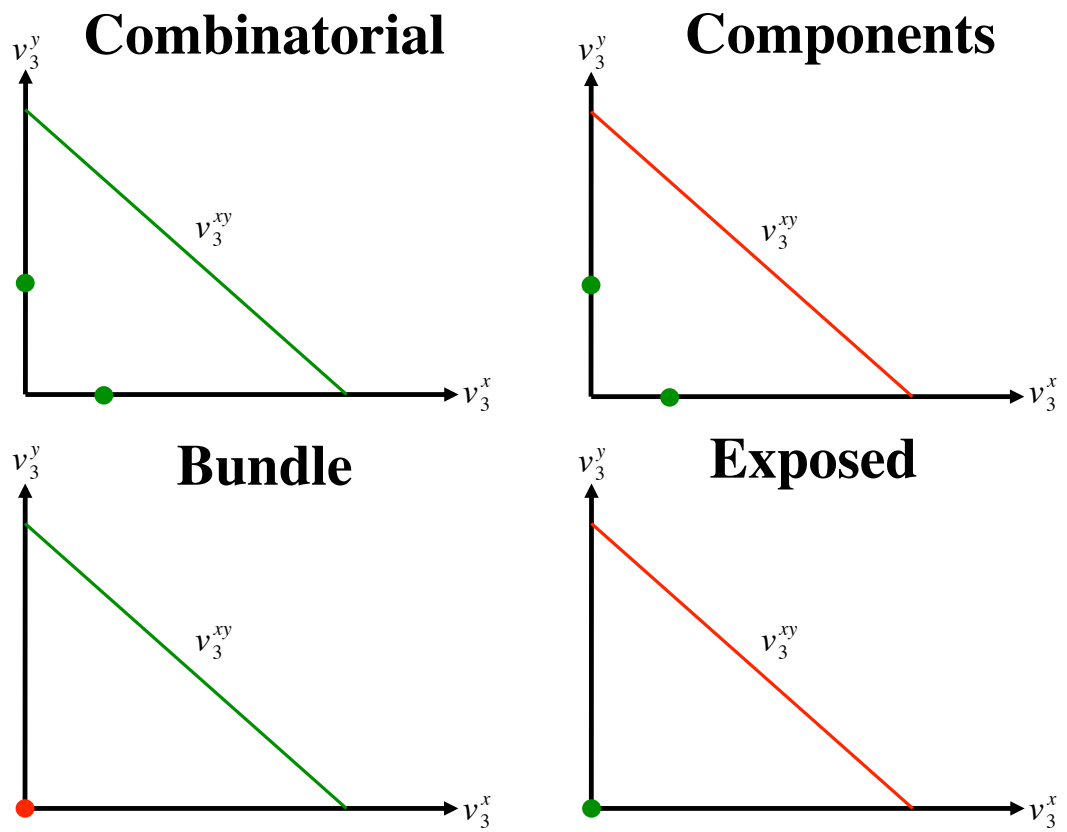
The Combinatorial treatment is a setting of a standard combinatorial auction where bidders can place bids on any subset of items.¹² The Exposed and Components treatments correspond to standard non-package auction settings, in which we simply vary global Bidder 3’s level of exposure. Figure 4.1 illustrates the treatments.

4.3.3 Auctions

In these experiments, we compare four sealed-bid auctions. The auctioneer takes the bidders’ bids and maps them into payments. Let us denote by $\mathbf{b} = (b_1^x, b_2^y, b_3^x, b_3^y, b_3^{xy}) \geq \mathbf{0}$ the vector of weakly positive bids. Bidding restrictions require that $(b_1^x, b_2^y, b_3^x, b_3^y, b_3^{xy}) \leq (v_1^x, v_2^y, v_3^{xy}, v_3^x, v_3^{xy})$. Hence, in the Combinatorial treatment, if $v_3^{xy} = 150$, Bidder 3 may bid up to 150 on item x , up to 150 on item y , and up to 150 on package xy . In

¹²If a bidder does not place a bid on a subset, it is equivalent to a bid of zero on that subset.

Figure 4.1: Treatments in the auction experiments



Point: valuation of individual items. Line: valuation of the package. Red: not permitted to bid. Green: permitted to bid.

some treatments, Bidder 3 is not allowed to bid on individual items or the package, so, without loss of generality, we can assume that in those cases the bids on those items equal to $-\infty$.

In every auction, the auctioneer must first determine who the winner of each item in the auction is. Since the seller values both items at zero, no item can be left unsold at positive price. While, in general, this is a computationally demanding integer programming problem, in our simple setting it is straightforward. Call a set of bids *feasible* if there is at most one bid on each item x and y . A bid on the package xy counts as a bid on both items. The winner determination problem reduces to finding a set of feasible bids, such that the sum of these bids in this set exceeds the sum of the feasible bids in any other set. Let vector $\mathbf{w} \in \{0, 1\}^5$ be the vector of winning bids, where 1 indicates a winning bid and 0 indicates a losing bid.

There are five possible cases, although not all cases apply to every treatment:

1. If $b_1^x + b_2^y \geq \max\{b_1^x, b_3^x\} + \max\{b_2^y, b_3^y\}$ and $b_1^x + b_2^y \geq b_3^{xy}$, then $\mathbf{w}_1 = (1, 1, 0, 0, 0)$ and local bidders win the items.
2. If $b_1^x + b_3^y \geq \max\{b_1^x, b_3^x\} + \max\{b_2^y, b_3^y\}$ and $b_1^x + b_3^y \geq b_3^{xy}$, then $\mathbf{w}_2 = (1, 0, 0, 1, 0)$; Bidder 1 wins item x and Bidder 3 wins item y .
3. If $b_3^x + b_2^y \geq \max\{b_1^x, b_3^x\} + \max\{b_2^y, b_3^y\}$ and $b_3^x + b_2^y \geq b_3^{xy}$, then $\mathbf{w}_3 = (0, 1, 1, 0, 0)$; Bidder 3 wins item x and Bidder 2 wins item y .
4. If $b_3^x + b_3^y \geq \max\{b_1^x, b_3^x\} + \max\{b_2^y, b_3^y\}$ and $b_3^x + b_3^y \geq b_3^{xy}$, then $\mathbf{w}_4 = (0, 0, 1, 1, 0)$; Bidder 3 wins both items with his individual item bids.
5. If $b_3^{xy} \geq \max\{b_1^x, b_3^x\} + \max\{b_2^y, b_3^y\}$ and $b_3^{xy} \geq b_3^x + b_3^y$, then $\mathbf{w}_5 = (0, 0, 0, 0, 1)$; Bidder 3 wins both items with his bid on the package xy .

The first three cases are entirely symmetric as there are two winning bidders and two

winning bids. In the Exposed and Components treatments, the auctioneer simply needs to find the highest bid on both items. In the Bundle treatment, the auctioneer checks whether the bid on the package (by Bidder 3) is greater or less than the sum of the individual bids (by the local bidders). Finally, in the Combinatorial treatment, we check whether the sum of the highest bids on the individual items exceeds the package bid. In this treatment, we take into account all of Bidder 3's bids.¹³ Hence, global bidder's own bids compete with one another. For example, if the global bidder submits a package and an individual item bid on x , then his bid on x together with Bidder 2's bid on y may outbid the package bid.

Once the winners are determined, auctioneer calculates the payments for every bidder according to the payment rules of a particular auction $P(\mathbf{b}) = (p_1, p_2, p_3)$. Define $W = \{\mathbf{w}_1, \mathbf{w}_2, \mathbf{w}_3, \mathbf{w}_4, \mathbf{w}_5\}$. Then, formally, the winners are determined by solving

$$\mathcal{W}(\mathbf{b}) \in \arg \max_{\mathbf{w} \in W} \mathbf{b} \cdot \mathbf{w} \quad (4.1)$$

and ties are broken randomly, so vector of winning bids is always unique. Hence, highest sum of feasible bids from \mathbf{b} is

$$R(\mathbf{b}) = \max_{\mathbf{w} \in W} \mathbf{b} \cdot \mathbf{w} \quad (4.2)$$

A bidder pays zero for any item he bids for and doesn't win. There is free disposal, so the auctioneer is not forced to sell both items.¹⁴

4.3.3.1 First-price auction

This is a classic auction in which the winners simply pay their bids. The payment function is therefore

¹³This is the also called OR bidding. We do not consider XOR bidding in which the global bidder would be able to win the package only with a bid on the package in the Combinatorial treatment (Chernomaz and Levin, 2012).

¹⁴Here we depart from Marszalec (2011), who did not impose free disposal.

$$P^{FP}(\mathbf{b}) = \begin{cases} (b_1^x, b_2^y, 0) & \text{in Case 1} \\ (b_1^x, 0, b_3^y) & \text{in Case 2} \\ (0, b_2^y, b_3^x) & \text{in Case 3} \\ (0, 0, b_3^x + b_3^y) & \text{in Case 4} \\ (0, 0, b_3^{xy}) & \text{in Case 5} \end{cases} \quad (4.3)$$

It is clear that in the first-price auction, it is never optimal for the bidders to bid their true value. In order to make positive profit, bidders must *shade* i.e bid below their value. Therefore, in general, the first-price auction will not be fully efficient especially when there are substantial bidder asymmetries (Maskin and Riley, 2000). In the Exposed and Components treatments, the global bidder effectively competes for the items in two separate markets. Hence, these treatments are equivalent to running two separate, simultaneous, first-price, sealed-bid auctions. In a framework with local bidders, who have perfectly correlated values over their respective items, and a slightly different valuation function for the global bidder, Chernomaz and Levin (2012) show that it is always optimal for the global bidder to submit zero bids on the individual items in the Combinatorial treatment. This is a striking result, however, we do not expect it to hold in our setting because of ex post bidder asymmetry.

4.3.3.2 Vickrey auction

In this auction, bidders pay the externality they exert on other bidders by placing their bids. Define $\mathbf{b}(-b_1^x)$ as vector \mathbf{b} where b_1^x takes value $-\infty$. Then, let us call $R(\mathbf{b}(-b_1^x))$ the *absent revenue* of b_1^x . The absent revenue equals to highest sum of feasible bids when bid b_1^x is removed from the auction.

$$P^{VA}(\mathbf{b}) = \begin{cases} (R(\mathbf{b}(-b_1^x)) - b_2^y, R(\mathbf{b}(-b_2^y)) - b_1^x, 0) & \text{in Case 1} \\ (R(\mathbf{b}(-b_1^x)) - b_3^y, 0, R(\mathbf{b}(-b_3^y)) - b_1^x) & \text{in Case 2} \\ (0, R(\mathbf{b}(-b_2^y)) - b_3^x, R(\mathbf{b}(-b_3^x)) - b_2^y) & \text{in Case 3} \\ (0, 0, R(\mathbf{b}(-b_3^x, -b_3^y))) & \text{in Case 4} \\ (0, 0, R(\mathbf{b}(-b_3^{xy}))) & \text{in Case 5} \end{cases} \quad (4.4)$$

In Exposed and Components treatment, the Vickrey payments are equivalent to payments in two separate second-price sealed bid auctions. Krishna and Rosenthal (1996) analyse this setting (with a different valuation function for the global bidder). They find (Theorem 1, p. 6) that in equilibrium local bidders will always report their true values and the global bidder's bidding function in continuously increasing in his value up to a threshold. For any value above that threshold, the global bidder simply bids the upper bound of the local bidders' value. While we do not expect this result to hold precisely, we can expect similar qualitative theoretical results in our setting.

In the Bundle treatment, it is a dominant strategy for all bidders to bid their valuation. While one may expect that the same would hold in the Combinatorial treatment, we have made some adjustments to the standard Vickrey auction there. Instead of paying the externality of his presence in the auction, the global bidder pays the externality of his *bids* in the auction.¹⁵ Therefore, we treat the global bidder's individual item bids as if they were made by a different party i.e. as shills. Since we allow the global bidder's individual item bids to affect his payment, he can do at least as well off as in Bundle treatment by bidding zero on the individual items and reporting his true valuation of the package.¹⁶

¹⁵In the standard Vickrey auction, if the global seller won both items, he would pay the sum of the losers' bids.

¹⁶It is still a dominant strategy for local bidders to bid their value. Global bidder can also be at least as well off in the Combinatorial treatment as in the Components treatment by bidding zero on the package and following the same bidding strategy on the individual items.

If the global bidder pursues this strategy in the Combinatorial treatment, our version of Vickrey auction may be inefficient. Consider the following set of valuations: $v_1^x = v_2^y = 80$ and $v_3^x = 90$, $v_3^y = 10$ so $v_3^{xy} = 150$. If the local bidders bid truthfully and the global bidder bids zero on the individual items and full value on the package, then the local bidders win. However, in an efficient allocation the global Bidder 3 should get item x and Bidder 2 should get item y .

4.3.3.3 Vickrey-Nearest Rule

The Vickrey-Nearest Rule (Day and Cramton, 2008, 2012) attempts to deal with the low revenue problem of the Vickrey auction (see Introduction and Example 1 below) by projecting the Vickrey auction payments to the closest point on the MRC. If b_1^x and b_2^y are the winning bids, we can define the *underbid* as $R(\mathbf{b}(-b_1^x, -b_2^y))$ i.e. the highest sum of feasible bids when all winning bids are removed from the auction (this is also the minimum-revenue line level). Hence, unless the sum of the Vickrey payments equals to the underbid in which case $P^{VNR}(\mathbf{b}) = P^{VA}(\mathbf{b})$, we have that:

$$P^{VNR}(\mathbf{b}) = \begin{cases} \left(\frac{1}{2}R(\mathbf{b}(-b_1^x, -b_2^y)) + \frac{1}{2}(R(\mathbf{b}(-b_1^x)) - b_2^y) - \frac{1}{2}(R(\mathbf{b}(-b_2^y)) - b_1^x), \right. \\ \left. \frac{1}{2}R(\mathbf{b}(-b_1^x, -b_2^y)) + \frac{1}{2}(R(\mathbf{b}(-b_2^y)) - b_1^x) - \frac{1}{2}(R(\mathbf{b}(-b_1^x)) - b_2^y), 0 \right) & \text{in Case 1} \\ \left(\frac{1}{2}R(\mathbf{b}(-b_1^x, -b_3^y)) + \frac{1}{2}(R(\mathbf{b}(-b_1^x)) - b_3^y) - \frac{1}{2}(R(\mathbf{b}(-b_3^y)) - b_1^x), 0, \right. \\ \left. \frac{1}{2}R(\mathbf{b}(-b_1^x, -b_3^y)) + \frac{1}{2}(R(\mathbf{b}(-b_3^y)) - b_1^x) - \frac{1}{2}(R(\mathbf{b}(-b_1^x)) - b_3^y) \right) & \text{in Case 2} \\ \left(0, \frac{1}{2}R(\mathbf{b}(-b_2^y, -b_3^x)) + \frac{1}{2}(R(\mathbf{b}(-b_2^y)) - b_3^x) - \frac{1}{2}(R(\mathbf{b}(-b_3^x)) - b_2^y), \right. \\ \left. \frac{1}{2}R(\mathbf{b}(-b_2^y, -b_3^x)) + \frac{1}{2}(R(\mathbf{b}(-b_3^x)) - b_2^y) - \frac{1}{2}(R(\mathbf{b}(-b_2^y)) - b_3^x) \right) & \text{in Case 3} \\ (0, 0, R(\mathbf{b}(-b_3^x, -b_3^y))) & \text{in Case 4} \\ (0, 0, R(\mathbf{b}(-b_3^{xy}))) & \text{in Case 5} \end{cases} \quad (4.5)$$

Whenever the global bidder wins both items, the payment is simply the sum of the second-highest bids or the package bid (whichever is higher). This is because the Vickrey payments coincide with the MRC (see Example 2 below). Hence, in Cases 4 and 5, given a particular bid vector, the payments are identical to the Vickrey auction. As in Vickrey auction, the global bidder can make himself at least as well off in the Combinatorial treatment as in the Bundle treatment by submitting zero bids on the individual items. However, because, in general, payments in his auction do not coincide with the Vickrey auction, local bidders have an incentive to shade their bids (Ausubel and Baranov, 2010, Goeree and Lien, 2012). Another way to see this point is to examine the payment rule and note that in Cases 1, 2 and 3, bidders' payments depend positively on their own bids.

4.3.3.4 Reference Rule

Since the Vickrey-Nearest Rule creates incentives for the local bidders to misreport their valuations, it is worth knowing to what extent these incentives can be reduced without sacrificing revenues on the MRC. Erdil and Klemperer (2010) tackled this question by introducing the Reference Rule, which further decouples local bidders' payments from their bids. In this auction, the auctioneer first determines the relative payments for the items. In our case and in line with Marszalec (2011), we set the payment for item x to be 3 times higher the payment for item y .¹⁷ While this goes somewhat against the equity considerations that motivated Erdil and Klemperer (2010), it allows us to make the Reference Rule payments substantially different from the Vickrey-Nearest Rule in most cases (see Examples below). Once the prices are determined and the bids are submitted, the auctioneer finds the reference point $(\frac{3}{4}R(\mathbf{b}(-b_1^x, -b_2^y)), \frac{1}{4}R(\mathbf{b}(-b_1^x, -b_2^y)))$ – a point on the minimum-revenue line such that the relative payments for the items correspond to

¹⁷In fact, Marszalec (2011) tested how the relative payments affect bidding in his experiments. He found that there are substantial differences in bidding compared to the Vickrey-Nearest Rule only when the relative payments are very different.

the predetermined ratio. If this point is on the MRC, these are the payments. If the reference point is not on the MRC, then the payments are the point on the MRC closest (in Euclidian distance) to the reference point. Erdil and Klemperer (2010) show that this payment rule reduces the sum of the marginal incentives to shade for local bidders compared to the Vickrey-Nearest Rule. Define $M(b_i, b_j, b_k) = \min\{b_i, R(\mathbf{b}(-b_i, -b_j)) - b_k\}$ to be the maximum payment along the underbid for the item that b_i is placed on, given b_j is the other winning bid and b_k is the losing bid on the item. Hence, unless the sum of the Vickrey payments equals to the underbid in which case $P^{RR}(\mathbf{b}) = P^{VA}(\mathbf{b})$, payments in the Reference Rule are given in (4.6).

Several points are worth mentioning. First, is that the sum of the Vickrey-Nearest Rule and the Reference Rule payments always coincides and equals to the underbid. The only difference between these rules is the distribution of payments for the winners of individual items. Second, theoretically, in the Bundle and the Combinatorial treatments, the Reference Rule (and Vickrey-Nearest Rule) payments should only affect bidding on individual items. Whenever the global bidder wins both items, his payments in the Vickrey-Nearest Rule, the Reference Rule, and the Vickrey auction coincide. Third, in the Exposed and Components treatments, the incentives of all bidders in the Vickrey, Vickrey-Nearest Rule and the Reference Rule auctions are identical. Fourth, analytical solutions for the optimal bidding strategies of local bidders for the Reference Rule do not exist even in the simplest setting of Ausubel and Baranov (2010). Therefore, in our experiments, we are interested in testing whether the Reference Rule does indeed lower local bidders' incentives to shade compared to the Vickrey-Nearest Rule.

$$P^{RR}(\mathbf{b}) = \left\{ \begin{array}{l} \text{Case 1} \left\{ \begin{array}{ll} \left(\frac{3}{4}R(\mathbf{b}(-b_1^x, -b_2^y)), \frac{1}{4}R(\mathbf{b}(-b_1^x, -b_2^y)), 0 \right) & \text{if } \frac{3}{4}R(\mathbf{b}(-b_1^x, -b_2^y)) \leq M(b_1^x, b_2^y, b_3^y); \frac{1}{4}R(\mathbf{b}(-b_1^x, -b_2^y)) \leq M(b_2^y, b_1^x, b_3^x) \\ \left(M(b_1^x, b_2^y, b_3^y), R(\mathbf{b}(-b_1^x, -b_2^y)) - M(b_1^x, b_2^y, b_3^y), 0 \right) & \text{if } \frac{3}{4}R(\mathbf{b}(-b_1^x, -b_2^y)) > M(b_1^x, b_2^y, b_3^y); \frac{1}{4}R(\mathbf{b}(-b_1^x, -b_2^y)) \leq M(b_2^y, b_1^x, b_3^x) \\ \left(R(\mathbf{b}(-b_1^x, -b_2^y)) - M(b_2^y, b_1^x, b_3^x), M(b_2^y, b_1^x, b_3^x), 0 \right) & \text{if } \frac{3}{4}R(\mathbf{b}(-b_1^x, -b_2^y)) \leq M(b_1^x, b_2^y, b_3^y); \frac{1}{4}R(\mathbf{b}(-b_1^x, -b_2^y)) > M(b_2^y, b_1^x, b_3^x); \end{array} \right. \\ \text{Case 2} \left\{ \begin{array}{ll} \left(\frac{3}{4}R(\mathbf{b}(-b_1^x, -b_3^y)), 0, \frac{1}{4}R(\mathbf{b}(-b_1^x, -b_3^y)) \right) & \text{if } \frac{3}{4}R(\mathbf{b}(-b_1^x, -b_3^y)) \leq M(b_1^x, b_3^y, b_2^y); \frac{1}{4}R(\mathbf{b}(-b_1^x, -b_3^y)) \leq M(b_3^y, b_1^x, b_3^x) \\ \left(M(b_1^x, b_3^y, b_2^y), 0, R(\mathbf{b}(-b_1^x, -b_3^y)) - M(b_1^x, b_3^y, b_2^y) \right) & \text{if } \frac{3}{4}R(\mathbf{b}(-b_1^x, -b_3^y)) > M(b_1^x, b_3^y, b_2^y); \frac{1}{4}R(\mathbf{b}(-b_1^x, -b_3^y)) \leq M(b_3^y, b_1^x, b_3^x) \\ \left(R(\mathbf{b}(-b_1^x, -b_3^y)) - M(b_3^y, b_1^x, b_3^x), 0, M(b_3^y, b_1^x, b_3^x) \right) & \text{if } \frac{3}{4}R(\mathbf{b}(-b_1^x, -b_3^y)) \leq M(b_1^x, b_3^y, b_2^y); \frac{1}{4}R(\mathbf{b}(-b_1^x, -b_3^y)) > M(b_3^y, b_1^x, b_3^x) \end{array} \right. \\ \text{Case 3} \left\{ \begin{array}{ll} \left(0, \frac{1}{4}R(\mathbf{b}(-b_3^x, -b_2^y)), \frac{3}{4}R(\mathbf{b}(-b_3^x, -b_2^y)) \right) & \text{if } \frac{3}{4}R(\mathbf{b}(-b_3^x, -b_2^y)) \leq M(b_3^x, b_2^y, b_3^y); \frac{1}{4}R(\mathbf{b}(-b_3^x, -b_2^y)) \leq M(b_2^y, b_3^x, b_1^x) \\ \left(0, R(\mathbf{b}(-b_3^x, -b_2^y)) - M(b_3^x, b_2^y, b_3^y), M(b_3^x, b_2^y, b_3^y) \right) & \text{if } \frac{3}{4}R(\mathbf{b}(-b_3^x, -b_2^y)) > M(b_3^x, b_2^y, b_3^y); \frac{1}{4}R(\mathbf{b}(-b_3^x, -b_2^y)) \leq M(b_2^y, b_3^x, b_1^x) \\ \left(0, M(b_2^y, b_3^x, b_1^x), R(\mathbf{b}(-b_3^x, -b_2^y)) - M(b_2^y, b_3^x, b_1^x) \right) & \text{if } \frac{3}{4}R(\mathbf{b}(-b_3^x, -b_2^y)) \leq M(b_3^x, b_2^y, b_3^y); \frac{1}{4}R(\mathbf{b}(-b_3^x, -b_2^y)) > M(b_2^y, b_3^x, b_1^x) \end{array} \right. \\ \text{Case 4} \left\{ (0, 0, R(\mathbf{b}(-b_3^x, -b_3^y))) \right. \\ \text{Case 5} \left\{ (0, 0, R(\mathbf{b}(-b_3^{xy}))) \right. \end{array} \right.$$

(4.6)

4.3.3.5 Illustration of all four auctions

The following two examples illustrate how the payments are calculated in all four auctions.

Example 1 Let us illustrate the auction payment rules by considering a bidding vector $\mathbf{b} = (48, 56, -\infty, -\infty, 60)$ in the Bundle treatment. We are in Case 1, because $b_1^x + b_2^y = 104 > 60 = b_3^{xy}$, so the two local bidders win the items. In the first-price auction, they simply pay their bids so $P^{FP}(\mathbf{b}) = (48, 56, 0)$. For the Vickrey auction, we first calculate the absent revenue and find that $R(\mathbf{b}(-b_1^x)) = R(\mathbf{b}(-b_2^y)) = 60$ so the highest sum of feasible bids is 60 whenever one of the winning bids is removed. Hence,

$$\begin{aligned}
 P^{VA}(\mathbf{b}) &= (R(\mathbf{b}(-b_1^x)) - b_2^y, R(\mathbf{b}(-b_2^y)) - b_1^x, 0) & (4.7) \\
 &= (60 - 56, 60 - 48, 0) \\
 &= (4, 12, 0)
 \end{aligned}$$

The Vickrey revenue of 16 is low and the Vickrey payments lie outside the core. The Vickrey-Nearest Rule projects the Vickrey payments directly on to the minimum-

revenue core, $R(\mathbf{b}(-b_1^x, -b_2^y)) = 60$, yielding

$$\begin{aligned}
P^{VNR}(\mathbf{b}) &= \left(\frac{1}{2}R(\mathbf{b}(-b_1^x, -b_2^y)) + \frac{1}{2}(R(\mathbf{b}(-b_1^x)) - b_2^y) - \frac{1}{2}(R(\mathbf{b}(-b_2^y)) - b_1^x), \right. \\
&\quad \left. \frac{1}{2}R(\mathbf{b}(-b_1^x, -b_2^y)) + \frac{1}{2}(R(\mathbf{b}(-b_2^y)) - b_1^x) - \frac{1}{2}(R(\mathbf{b}(-b_1^x)) - b_2^y), 0 \right) \\
&= \left(\frac{1}{2} \times 60 + \frac{1}{2} \times 4 - \frac{1}{2} \times 12, \frac{1}{2} \times 60 + \frac{1}{2} \times 12 - \frac{1}{2} \times 4 \right) \\
&= (26, 34, 0)
\end{aligned} \tag{4.8}$$

Finally, in the Reference Rule auction, we first calculate the reference point

$$\left(\frac{3}{4}R(\mathbf{b}(-b_1^x, -b_2^y)), \frac{1}{4}R(\mathbf{b}(-b_1^x, -b_2^y)) \right) = (45, 15) \tag{4.9}$$

and find that, since $\frac{3}{4}R(\mathbf{b}(-b_1^x, -b_2^y)) = 45 < 48 = b_1^x$ and $\frac{1}{4}R(\mathbf{b}(-b_1^x, -b_2^y)) = 15 < 56 = b_1^y$,¹⁸ the payments are indeed equal to the reference point

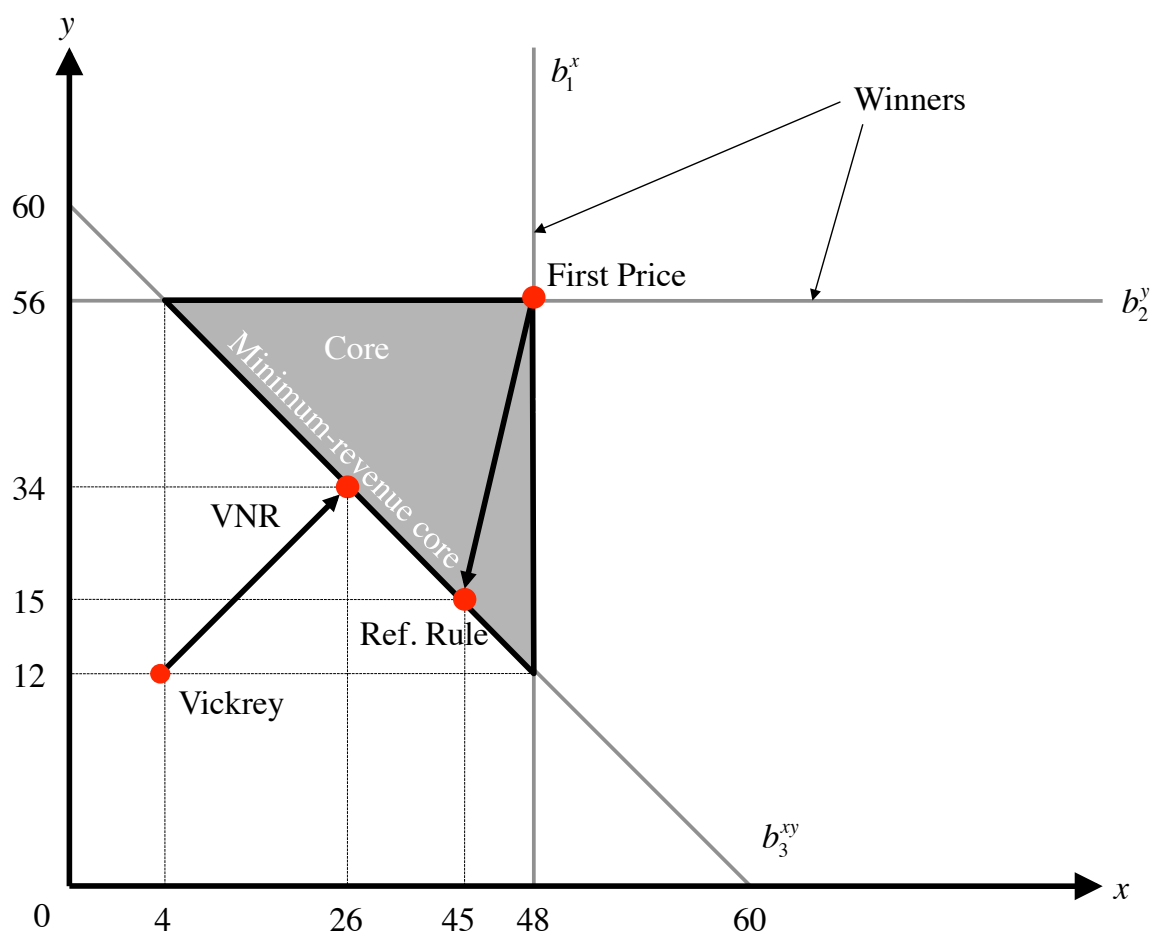
$$\begin{aligned}
P^{RR}(\mathbf{b}) &= \left(\frac{3}{4}R(\mathbf{b}(-b_1^x, -b_2^y)), \frac{1}{4}R(\mathbf{b}(-b_1^x, -b_2^y)), 0 \right) \\
&= (45, 15, 0)
\end{aligned} \tag{4.10}$$

Figure 4.2 illustrates the bids and the payments for all four rules in Example 1. ■

Example 2 In this example, let us consider a bidding vector $\mathbf{b} = (20, 50, 46, 22, 160)$ in the Combinatorial treatment. We are in Case 5, because Bidder 3's package bid of 160 exceeds the sum of any other feasible set of bids, so Bidder 3 wins both items with his package bid. In the first-price auction, Bidder 3 simply pays his bid so $P^{FP}(\mathbf{b}) = (0, 0, 160)$. In the other three auction rules, Bidder 3's payments coincide with the

¹⁸ $M(b_1^x, b_2^y, b_3^y) = \min\{48, 60 + \infty\} = 48$ and $M(b_1^x, b_2^y, b_3^y) = \min\{56, 60 + \infty\} = 56$

Figure 4.2: Illustration of all four rules in Case 1 (Example 1)



minimum-revenue core, so

$$P^{VA}(\mathbf{b}) = P^{VNR}(\mathbf{b}) = P^{RR}(\mathbf{b}) = (0, 0, R(\mathbf{b}(-b_3^{xy}))) = (0, 0, 96) \quad (4.11)$$

Note that the payment of Bidder 3 is affected by his bid on item y . Bidder 3 could have himself strictly better off in the Vickrey and both MRCS auctions by submitting zero bids on both individual items. Figure 4.3 illustrates the bids and the payments for all four rules in Example 2. ■

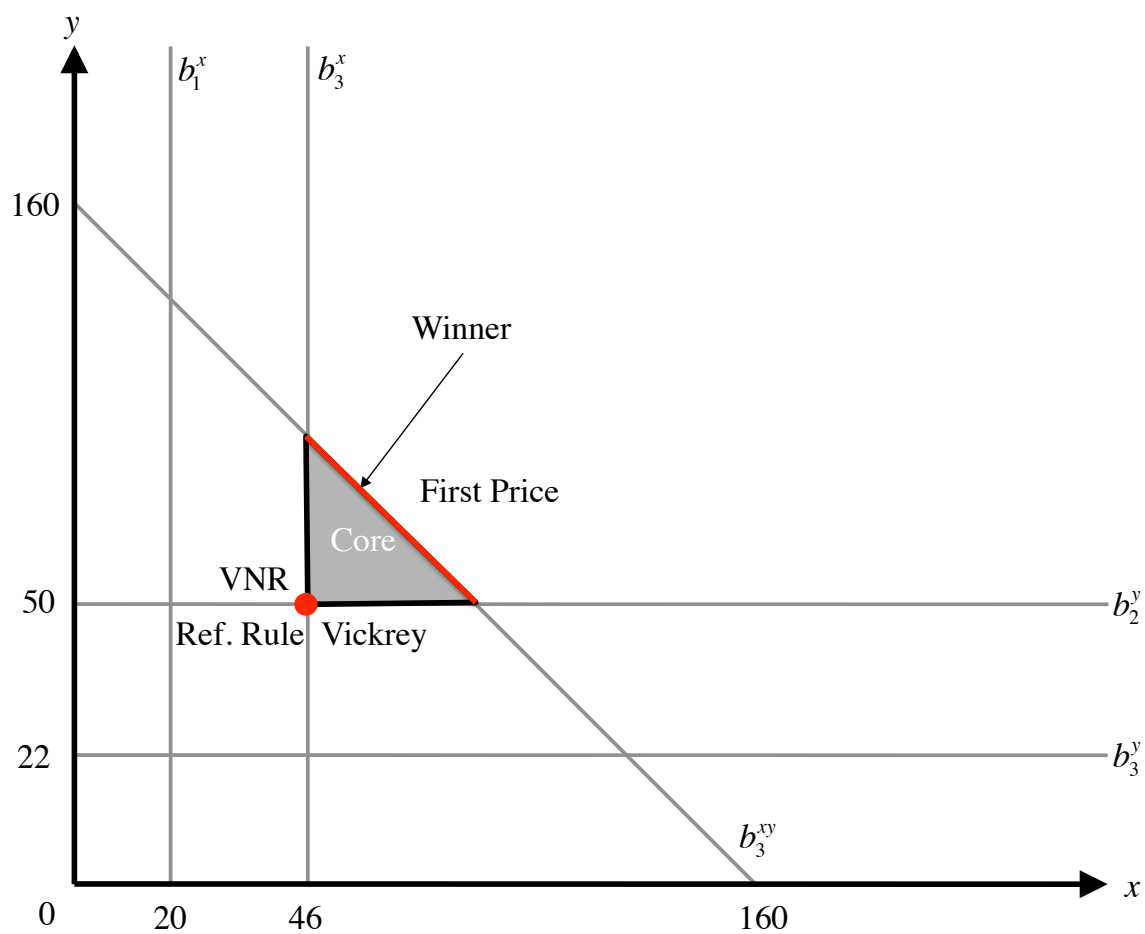
4.3.4 Hypotheses

Although no full analytic solutions are known for any of the auctions with the valuation functions we proposed, we can still rely on related or partial solutions as well as other experimental evidence to test several hypotheses in our data.

A. First-price auction dominates the Vickrey auction in terms of revenue with MRCS auctions in between. The experimental evidence of Marszalec (2011) strongly rejected the theoretical predictions of Ausubel and Baranov (2010) that in equilibrium (albeit in a simple setting), the Vickrey auction should dominate the first-price auction in terms of efficiency and revenue. The revenue ranking was the following: first-price, followed by MRCS auctions, followed by the Vickrey auction. In another series of experiments, Marszalec and I have found that this conclusion is robust under a variety of competition and information treatments. Here, we test whether these results are robust to different valuation and package treatments.

B. Exposure lowers global bidders' profits. The Exposed and the Components treatments allow us to test directly the extent to which the global bidder is affected by exposure. Chernomaz and Levin (2012) fix the valuation of the individual items and vary the strength of complementarities (parameter β). We do the opposite by fixing the distribution of the valuation of the package and varying the valuation of individual

Figure 4.3: Illustration of all four rules in Case 5 (Example 2)



items. This allows us to test directly for exposure rather than for the effect of strength of package complementarities.

C. Global bidders are at least as well off under the Combinatorial treatment as under the Bundle/Components treatments in all auctions. The Combinatorial treatment gives the global bidder an opportunity to bid on individual items. In the descriptions of the auctions, we explained that the global bidder can leave himself at least as well off as in the Bundle treatment by simply bidding zero on the individual items and submitting a positive bid on the package.

D. Package bidding creates threshold effects in the Bundle treatment. In the Bundle treatment, local bidders face a coordination problem in the first-price, Vickrey-Nearest Rule, and Reference Rule auctions when trying to outbid the global bidder. In the Exposed and Components treatments (and in the Vickrey auction), there is no coordination problem since bidders compete in separate markets.

E. Local bidders bid their value in the Vickrey auction. It is a dominant strategy for the local bidders to bid their value in all treatments of the Vickrey auction. However, the particular design of the Vickrey auction sometimes has an effect on the extent to which it is indeed demand-revealing (Knetsch et al., 2001). Our setting is reasonably transparent for local bidders in the Vickrey auction, however, there is no reason to assume ex ante that local bidders will behave according to their dominant strategy.

F. Local bidders shade less in the Reference Rule than in the Vickrey - Nearest Rule. While these two auctions must coincide in terms of revenue, Erdil and Klemperer (2010) showed that the Reference Rule gives local bidders lower marginal incentives to shade in a simple setting. Marszalec (2011) did not find any evidence for this, except in treatments where the relative payments for the items were very different.

We can stress-test whether the prediction of Erdil and Klemperer (2010) holds in our Bundle and Combinatorial treatments.

4.3.4.1 Statistical tests and methodology

In line with a lot experimental literature, we use two robust non-parametric tests to test all our hypotheses. For bidding behaviour (including shading), we use the Mann-Whitney tests for stochastic dominance. This test allows us to say when a treatment or an auction induces more aggressive bidding behaviour. Since the variable outcomes in our sub-populations are unlikely to be normal, the Mann-Whitney and Kruskal-Wallis tests are preferred to the conventional t - and F -tests.¹⁹ In order to test auction outcomes, such as revenue and allocative efficiency, we use the Hodges-Lehmann median-difference test. The null hypothesis for this test is that the median difference between values of two sub-populations is zero. This test dispenses with the usual assumptions of equal variability and normality in the subpopulations and allows us to compare the homogeneity of distributions more robustly. Finally, since we are potentially comparing up to four treatments to one another and testing up to six null hypotheses simultaneously, we use the Holm–Bonferroni method in order to control the familywise error rate.

4.4 Experimental procedures

The experiments were conducted at the Centre for Experimental Social Science at Nuffield College, University of Oxford. We ran 9 sessions in December 2012-January 2013. We targeted science, technology, engineering, mathematics (STEM) and medical students in order to make sure that participants are comfortable with elementary algebra. Sessions lasted between 2 and 3 hours.²⁰ There were 174 participants in total

¹⁹Mann-Whitney test is more efficient than the t -test when sub-populations are not normally distributed.

²⁰Nevertheless, students reported a high level of satisfaction with the experiments. There were at least two breaks in each session during which subjects were invited to drink some water, have a sugary

(numbers varied between 12 and 25 per session). A short, voluntary participant questionnaire was handed out after each session. Participant characteristics are described in Table 4.10 of Appendix 4.7.1.

Instructions with examples were read out to all the participants. They were then asked to complete a test.²¹ If they made a mistake, a supervisor would indicate that a mistake has been made. If the mistake was serious enough to indicate a lack of understanding of the auction rules (i.e. not just an obvious calculation error), the participants would be told they had not passed the test and could not participate in the experiment. Participants had at most half an hour to complete the test.²² After the subjects passed the test, they remained in the lab for the rest of the experimental session.

Before each of the four treatments, the supervisor read out three examples illustrating the treatment. Participants played two practice rounds, followed by ten payoff-relevant rounds. In every round, each participant was anonymously and randomly matched with two other participants. Bidder roles (i.e. Bidder 1, 2 or 3) were assigned randomly, but participants were told that there would be two local bidders and one global bidder in each round. Participant's valuation for the item(s) was randomly assigned in every round according to the valuation functions and the treatments described above. Figure 4.5a in Appendix 4.7.1.1 gives an illustration of a global bidder's screen in the Combinatorial treatment. After every participant submitted a bid in a round, all participants were told the outcome of the round. Participants were given extensive feedback (see Figure 4.5b in Appendix 4.7.1.1). For a particular auction, treatments were done in a different order in different sessions to minimise learning effects. The software was coded in z-Tree (Fischbacher, 2007) and all calculations were checked in

snack or go to the bathroom. They were not permitted to talk.

²¹See Appendix 4.7.3 for an example of the Vickrey auctions instructions and test.

²²The test was discriminating. In sessions 7, 8, and 9, we did not target STEM and medical students on the mailing list. In these sessions, the pass rate was much lower. See Table 4.11 in Appendix 4.7.1.

R.

Participants were paid anonymously according to their profit in two randomly selected rounds of each treatment. We did not pay for every round in order to avoid income effects. Participant payments ranged between £12 and £60 with an average of £40. On a per hour basis, this average corresponds roughly to a research assistant's wage at the University of Oxford.

4.5 Results

4.5.1 Baseline results

We now present the results of our experiments. Our dataset has 160 rounds per treatment for the Vickrey-Nearest Rule and 140 rounds per treatment for every other auction. We first give an overview of the baseline results and then systematically test all the hypotheses. Table 4.1 summarises revenue and efficiency across every auction and treatment. We calculate efficiency in every round in the usual way

$$\text{efficiency} = \frac{\text{sum of winners' valuations}}{\text{optimal sum of bidders' valuations}} \times 100$$

In most rounds, items are allocated to bidders who jointly value the items the most. The median efficiency for every auction and treatment is 100, so we do not report it in the tables. The results in this table suggest that the first-price auction revenue-dominates all other auctions and we test this hypothesis below.

Table 4.2 gives a summary of local and global bidders' average profits. Bidder surplus is defined as the sum of winners' profits, so the local and global bidders' profits add up to the total surplus.

Table 4.1: Revenue and efficiency

Auction	Variable	Treatment			
		Exposed	Components	Bundle	Combinatorial
First-price N=560	Revenue	125.8 135.8 (37.4)	120.7 124.5 (32.4)	119.3 121.8 (28.3)	121.6 129.0 (29.4)
	Efficiency	86.3 (22.1)	92.8 (11.7)	98.7 (5.15)	98.9 (3.42)
Vickrey N=560	Revenue	76.4 82.5 (48.5)	78.6 79.5 (37.9)	80.3 80.5 (43.4)	98.9 98.0 (43.2)
	Efficiency	95.1 (12.2)	94.0 (12.5)	100.0 (0.07)	98.5 (5.42)
VNR N=640	Revenue	84.6 90.0 (42.2)	75.6 71.3 (37.4)	83.9 81.5 (33.9)	100.6 101.5 (43.5)
	Efficiency	94.0 (14.2)	94.6 (10.6)	99.4 (3.17)	96.7 (7.26)
Ref. Rule N=560	Revenue	81.4 90.5 (47.2)	75.2 74.5 (38.6)	86.0 89.0 (37.3)	95.7 95.0 (39.7)
	Efficiency	93.0 (14.7)	93.3 (13.3)	99.0 (7.50)	96.9 (6.93)

Comment: Revenue: mean | median; Efficiency: mean; standard deviations in parentheses

Table 4.2: Profits and bidder surplus

Auction	Variable	Treatment			
		Exposed	Components	Bundle	Combinatorial
First-price N=560	Local bidder profit	9.05 (12.8)	8.78 (11.3)	6.27 (13.0)	6.19 (10.5)
	Global bidder profit	-0.98 (40.6)	19.2 (28.5)	26.5 (23.1)	26.5 (22.2)
	Bidder surplus	8.1 (37.9)	28.0 (24.4)	32.8 (19.1)	32.7 (17.8)
Vickrey N=560	Local bidder profit	17.3 (40.8)	20.2 (31.0)	20.9 (41.3)	15.2 (33.3)
	Global bidder profit	49.6 (54.2)	48.5 (50.5)	52.7 (49.2)	43.8 (47.1)
	Bidder surplus	66.9 (52.8)	68.7 (44.3)	73.6 (43.7)	59.0 (45.5)
VNR N=640	Local bidder profit	16.3 (37.4)	21.7 (31.6)	13.4 (29.3)	15.6 (30.6)
	Global bidder profit	49.5 (55.6)	50.4 (55.3)	52.7 (48.6)	35.1 (43.8)
	Bidder surplus	65.8 (52.6)	72.1 (46.4)	66.1 (42.5)	50.7 (42.2)
Ref. Rule N=560	Local bidder profit	17.2 (33.5)	22.2 (30.4)	11.8 (28.8)	11.7 (27.6)
	Global bidder profit	47.9 (62.0)	51.2 (51.6)	58.6 (50.2)	41.0 (48.3)
	Bidder surplus	65.1 (54.0)	73.4 (45.2)	70.4 (44.3)	52.7 (45.6)

Comment: Means reported; standard deviations in parentheses

4.5.2 Testing the hypotheses

A. First-price auction dominates the Vickrey auction in terms of revenue with MRCS auctions in between. Table 4.3 shows that the first-price auction dominates all other auctions in terms of revenue in every treatment. The tests based on Hodges-Lehmann estimator reject the equality hypothesis at over 99% confidence level, even after applying the Holm-Bonferroni correction.²³ However, we cannot rank the other three auctions. Contrary to the main theoretical predictions in MRCS auction literature, we do not find any evidence that MRCS auctions produce more revenue than the Vickrey auction. These results are not surprising (especially in the Exposed and Components treatments where they are identical) and they are entirely in line with Marszalec (2011). In terms of magnitude, the additional revenue in the first-price auction almost doubles in Exposed and Components treatments compared to the Combinatorial treatment. The mirror-image of these results is that bidder surplus in the first-price auction is dominated by the other auctions. These differences in surplus are entirely driven by changes in global bidders' profits. We cannot reject the null hypothesis that local bidders' profits are the same across auctions in any treatment (see Tables in Appendix 4.7.2.1).

Unsurprisingly, we cannot reject the hypothesis of homogeneity of efficiency across the auctions in any treatment using median-difference tests. The presence of a large number of ties (rounds in which efficiency equals 100) creates difficulties for non-parametric tests. This becomes evident when we run Kruskal-Wallis tests adjusted and unadjusted for ties to test whether efficiency is homogeneous across all auctions for every treatment. In the Combinatorial and Exposed treatments, we obtain a clear rejection of equality from both statistics (all p -values < 0.065); in the Components treatment,

²³A careful reader will note that in these tables we reported the median differences between values of two sub-populations rather than the difference between the medians of two sub-populations. The latter results can be immediately obtained from Table 4.1.

Table 4.3: Median-difference test results — Revenue

Exposed	Vickrey [N=140]	First-price [N=140]	VNR [N=160]	Ref. Rule [N=140]
Vickrey [N=140]	—	—	—	—
First-price [N=140]	50.0*** [36.0:64.0]	—	—	—
VNR [N=160]	8.0 [0.0:17.0]	-42.0*** [-54.0:-31.0]	—	—
Ref. Rule [N=140]	5.0 [-3.0:15.0]	-45.0*** [-57.0:-32.0]	-2.0 [-11.0:6.0]	—
Components	Vickrey [N=140]	First-price [N=140]	VNR [N=160]	Ref. Rule [N=140]
Vickrey [N=140]	—	—	—	—
First-price [N=140]	44.0*** [32.0:56.0]	—	—	—
VNR [N=160]	-4.0 [-12.0:4.0]	-49.0*** [-60.0:-38.0]	—	—
Ref. Rule [N=140]	-4.0 [-12.0:4.0]	-49.0*** [-60.0:-37.0]	0.0 [-7.2:7.0]	—
Bundle	Vickrey [N=140]	First-price [N=140]	VNR [N=160]	Ref. Rule [N=140]
Vickrey [N=140]	—	—	—	—
First-price [N=140]	39.0*** [27.0:52.0]	—	—	—
VNR [N=160]	3.0 [-5.0:12.0]	-36.0*** [-46.0:-27.0]	—	—
Ref. Rule [N=140]	5.5 [-3.0:14.0]	-33.0*** [-44.0:-23.0]	3.0 [-4.0:10.0]	—
Combinatorial	Vickrey [N=140]	First-price [N=140]	VNR [N=160]	Ref. Rule [N=140]
Vickrey [N=140]	—	—	—	—
First-price [N=140]	25.0*** [14.0:36.0]	—	—	—
VNR [N=160]	2.9 [-6.0:11.0]	-21.0*** [-33.0:-9.0]	—	—
Ref. Rule [N=140]	-3.0 [-11.0:5.0]	-27.0*** [-39.0:-15.0]	-5.5 [-14.0:3.0]	—
Reported median difference row and column values				
<i>Comment:</i>	***/**/* denote rejection of H_0 at the 1/5/10% significance levels			
	[Confidence interval at the appropriate significance level (default 90%)]			
	No tests are affected by Holm-Bonferroni correction			

there is no evidence against equality (adj. p -value=0.660 and unadj. p -value=0.509). However, in the Bundle treatment, there is ambiguous evidence – adj. p -value=0.017 and unadj. p -value=0.725 – which is caused by the ties. Pairwise Mann-Whitney tests also appear to indicate that first-price auction is less efficient than the other auctions in the Exposed treatment. This exceptional sensitivity of Mann-Whitney and Kruskal-Wallis tests to adjustment for ties casts doubt over the validity of their conclusions. Therefore, we do not have any robust evidence against homogeneity of efficiency across auctions.

In order to estimate whether full efficiency occurs unusually frequently, we fit a flexible zero-inflated negative binomial (or Pólya) model to the data. Zero-inflated models are frequently used in estimation of datasets that are dispersed, non-normal, and heavily skewed towards zero.²⁴ In this model, an auction is fully efficient (i.e. efficiency loss is zero) with probability ϕ .²⁵ However, with probability $1 - \phi$, bidders shade sufficiently to cause efficiency loss and this loss is distributed according to the negative binomial distribution with parameter μ on the truncated interval $[0, 100]$. Hence

$$f(\text{efficiency loss} = k | \phi, \lambda, \alpha) = \begin{cases} \phi + (1 - \phi)(1 - \lambda\alpha)^{-\frac{1}{\alpha}} & k = 0 \\ (1 - \phi)(1 - \lambda\alpha)^{-\frac{1}{\alpha}} \frac{\Gamma(k + \frac{1}{\alpha})}{k! \Gamma(\frac{1}{\alpha})} (1 + \frac{1}{\lambda\alpha})^{-k} & k > 0 \end{cases}$$

where $0 \leq \phi \leq 1$, $\lambda > 0$, Γ is the gamma function, and $\alpha > 0$ is the “dispersion” parameter. If $\phi = 0$, we obtain the standard negative binomial model and as α tends to zero, the distribution approaches the zero-inflated Poisson distribution (with non-inflated mean and variance λ). This model resembles the “spike-logit” model developed by Costa-Gomes and Crawford (2006) to identify exact type-specific responses in two-person guessing games. The logit error structure in their model implies that the dis-

²⁴In our setting, this is caused neither by sample selection nor by truncation so Heckman and Tobit models are not applicable.

²⁵Efficiency loss = 100 – realised efficiency.

tribution of errors was symmetric. However, in our setting, efficiency loss is restricted to be at most 0, therefore the distribution of efficiency loss is asymmetric and the logit model is inappropriate. However, the negative binomial distribution is asymmetric over a reasonable parameter range and flexible enough to allow us to be agnostic about the precise shape of the distribution of efficiency loss in the data.²⁶

The zero-inflated negative binomial distribution has a well-defined log-likelihood function, which we can maximise using standard optimisation software such as R and Stata. Table 4.4 reports the estimated coefficients. We are primarily interested in seeing whether the non-inflated model in which $\phi = 0$ performs better than the zero-inflated model. We test this by using the classical Vuong (1989) likelihood ratio test for nested models. In addition, we test whether α is zero i.e. if our data are better approximated by a zero-inflated Poisson model.

In the table, we report $\phi^* = \frac{1}{1+e^\phi}$ i.e. the transformed probability of full efficiency in each treatment as well as the estimate of the parameter λ . We find that the Vuong test rejects the null hypothesis of the absence of a spike at zero for the negative binomial model in 11 out of 16 treatments. The reason why the test did not reject the null in the five remaining treatments is almost certainly due to the fact that the non-zero sample was *too small* i.e. the spike at full efficiency was overwhelming in the data. In this case, the maximum likelihood estimations and the Vuong test can perform poorly (Shi, 2013). Finally, in every treatment, likelihood ratio (LR) tests indicate that the zero-inflated negative binomial model performs better than the zero-inflated Poisson model. This is not surprising since our efficiency data is dispersed (variance \gg mean) in every treatment and the Poisson model is suitable for data generating processes where the mean is roughly equal to the variance.

²⁶Of course, negative binomial is a count model and is therefore only appropriate in integer-valued data. Nevertheless, the log-likelihood function is expressed entirely in terms of the gamma function, which is continuous on $(0, \infty)$.

Table 4.4: Testing spikes in efficiency

Auction	Treatment	λ	ϕ^*	Vuong test	α	LR test for $\alpha = 0$
First-price	Exposed	3.685*** (0.86)	0.65*** (0.178)	4.53***	0.331 (0.075)	342.94***
	Components	2.988*** (0.833)	0.554*** (0.176)	5.16***	0.303 (0.072)	174.79***
	Bundle	2.730*** (0.182)	2.363*** (0.302)	1.81**	0.328 (0.167)	35.64***
	Combinatorial	2.125*** (0.152)	1.898*** (0.254)	2.06**	0.283 (0.149)	18.57***
Vickrey	Exposed	3.190*** (0.151)	1.373*** (0.212)	2.16**	0.591 (0.196)	211.75***
	Components	3.038*** (0.103)	0.912*** (0.187)	3.71***	0.375 (0.100)	218.51***
	Bundle [†]	—	—	—	—	—
	Combinatorial	2.524*** (0.245)	1.987*** (0.276)	— [‡]	0.820 (0.430)	80.82***
VNR	Exposed	3.392*** (0.115)	1.384*** (0.198)	2.47***	0.389 (0.107)	233.22***
	Components	2.837*** (0.105)	0.773*** (0.172)	5.36***	0.489 (0.120)	263.55***
	Bundle	2.537*** (0.260)	3.074*** (0.387)	1.09	0.386 (0.264)	20.20***
	Combinatorial	2.600*** (0.128)	1.14*** (0.188)	3.03***	0.530 (0.168)	126.67***
Ref. Rule	Exposed	3.360*** (0.116)	1.134*** (0.197)	2.82***	0.417 (0.116)	222.29***
	Components	3.005*** (0.114)	0.699*** (0.181)	1.24	0.546 (0.134)	349.28***
	Bundle	3.042*** (0.601)	2.954*** (0.502)	— [‡]	1.836 (1.879)	128.84***
	Combinatorial	2.404*** (0.176)	0.959*** (0.213)	2.66***	0.914 (0.362)	151.95***

Comments: Standard errors in parentheses. ***/**/* denote rejection of the appropriate H_0 at the 1/5/10% significance levels.

Vuong test for the zero-inflated vs. standard negative binomial model using z -score. LR-test of $\alpha = 0$ using $\chi^2(1)$.

[†]Vickrey Bundle data did not converge because there was only one non-zero observation.

[‡]One-sided test could not be performed because the non-inflated model performed better.

Table 4.5: Effect of exposure on global bidders' profit

	Profit in Exposed – Profit in Components			
	Conditional on winning		Unconditional	
	Median difference	<i>p</i> -value	Median difference	<i>p</i> -value
First-price	-18.0	0.002	-12.0	0.000

Comment: Reported *p*-value using the Somers' *D* statistic
Median difference between the values in the Exposed and the values in the Components treatment

B. Exposure lowers global bidders' profits. We now test to what extent exposure harms global bidders. While on average global bidders' valuation of the package between the Exposed and Components treatments is the same, the valuations of the individual items are drastically different. In the Exposed treatment, in which winning one item necessarily means negative profit, the global bidder faces much higher exposure in every auction compared to the Components treatment. We performed the median-difference tests on global bidders' profits between the Exposed and the Components treatments. We only find significant differences in global bidders' profits in the first-price auction, so we omit the other results from Table 4.5.

The reason for this is that global bidders appear to bid significantly more aggressively in the Exposed than in the Components treatment of Vickrey and MRCS auctions, but similarly in the first-price auction (see Table 4.6). However, local bidders bid similarly across these two treatments in all auctions.²⁷ Let's look at their median bids in the Vickrey auction. They are 101 on each item. This is not a coincidence because by bidding 101 the global bidders always outbid the local bidders on the item. However, that does not mean that this bid is always profitable. If a global bidder's valuation of the package is 150 and local bidders always bid their value, then 1 in 8 times the global bidder will make a loss.

To get a better sense of how global bidders respond to exposure, we estimate their

²⁷There is one exception. We can reject the null hypothesis that bidding is the same for the local bidders on item *x* (but not on item *y*) in the first-price auction, though only at a 10% significance level. Local bidders' bidding does appear slightly less aggressive in the Exposed treatment as one might expect.

Table 4.6: Global bidders’ bids on individual items

	Median bid					
	On item x			On item y		
	Exposed	Components	p -value	Exposed	Components	p -value
First-price	60.5	51.0	0.168	60.0	60.0	0.711
Vickrey	101	76.8	0.000	101	80	0.000
VNR	90.0	70.0	0.001	90.0	68.8	0.000
Ref. Rule	99.0	70.0	0.000	97.5	67.8	0.000

Comment:

Reported p -value using the Mann-Whitney U statistic

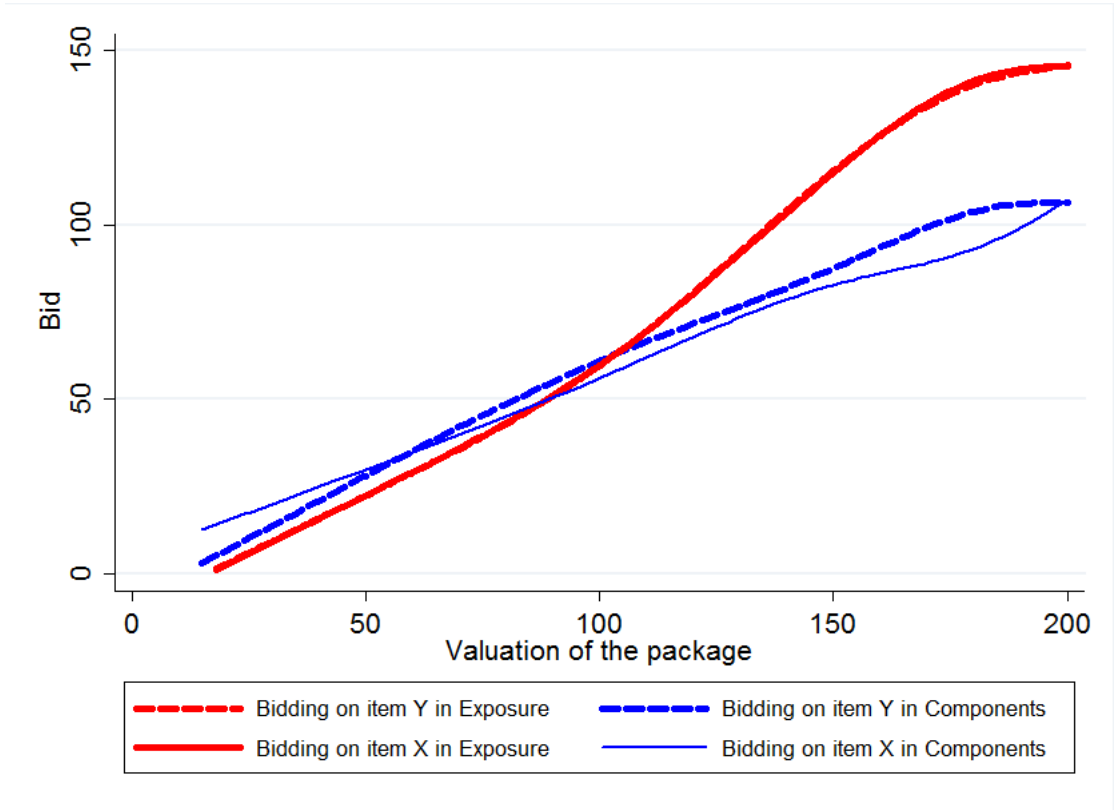
bidding functions using cubic splines. We divide the support of the valuation into six pieces and fit a smooth cubic polynomial function to their bids on each piece. In Figure 4.4, we illustrate the shift in the bidding function for the global bidders in the Vickrey auction (MRCS auctions are similar). Global bidders are clearly much more aggressive on both items (the red lines almost coincide) for valuations above 100 in the Exposed treatment. Our data indicate that global bidders act according to a “threshold” bidding function, similar to the one described by Krishna and Rosenthal (1996).²⁸

C. Global bidders are at least as well off under the Combinatorial treatment as under the Bundle/Components treatments in all auctions. This hypothesis tests whether global bidders can take into account the effect of their own bids on their payments. In the Combinatorial treatment, we gave global bidders an opportunity to bid on individual items as well as the package. Since this expands their strategy space compared to the Bundle or Components treatments, they should be made no worse off in the Combinatorial treatment.

However, Table 4.7 shows that this is not the the case. The global bidders raised lower profits, conditional on winning, in the Combinatorial treatment than in the Bundle treatment in both MRCS and Vickrey auctions. The effect is not significant in the first-price auction. In the Reference Rule auction, global bidders were better off in both the Components and the Bundle treatment than in the Combinatorial treatment.

²⁸Although the threshold bidding function of our bidders is almost certainly not optimal.

Figure 4.4: Estimated bidding functions for the global bidder in the Exposed and Components treatments of the Vickrey auction



The profit loss does not disappear when we do not condition on winning in the MRCS auctions, so overall global bidders do not simply increase their probabilities of winning. Revenues for the auctioneer were significantly higher in the Combinatorial treatment than in other treatments in both MRCS and Vickrey auctions, but not in the first-price auction (see Tables in Appendix 4.7.2.2). Therefore, we cannot explain higher revenues simply by increased competition on the individual items.

The reason why bidders hurt themselves in the Combinatorial treatment is because their bids affect their own payments in the MRCS and Vickrey auctions. We ran a counterfactual calculation to see how much profit bidders lost by affecting their payments. In the Vickrey and the MRCS auctions, we set the global bidders' non-winning bids to zero. In the first-price auction, we recorded whether global bidders could have made more profit by dropping one of their bids. Hence, if the global bidder won the item with his package bid, but could have won it with the individual item bids, we recorded the opportunity cost of the package bid. The results of these counterfactual calculations are reported in Table 4.8. Winning global bidders lose approximately 20 per cent of their profits because of their other bids. We call this inability to take into account the effect of your own bids on your payments the *combinatorial curse*.

As we mentioned in the description of the treatments, this setting can be interpreted as a setting in which the global bidder attempts to shill by having the shilling agent bid on the individual items. However, this experimental setting underlines the fact that bidder shilling is unlikely whenever there is competition on the items. Opportunities to shill for the global bidder led to lower profits.²⁹

D. Package bidding creates threshold effects in the Bundle treatment. An ideal way to test for more aggressive bidding behaviour is to look at shading. However,

²⁹In dynamic combinatorial clock auctions used in practice, the auctioneer uses all the past bids to calculate excess demand in every round. Therefore, all past bids of the bidders may affect final payments.

Table 4.7: Median-difference test results — Global bidders' profit conditional on winning

	Components	Bundle
First-price	[N=116]	[N=104]
Combinatorial	7.5**	-4.0
[N=117]	[1.0:14.5]	[-8.5:0.5]
Vickrey	[N=114]	[N=103]
Combinatorial	-6.5	-18.5***
[N=114]	[-17.5:4.0]	[-35.0:-2.0]
VNR	[N=138]	[N=120]
Combinatorial	-11.5*	-26.0***
[N=124]	[-22.0:-1.0]	[-42.0:-10.5]
Ref. Rule	[N=112]	[N=110]
Combinatorial	-18.4***	-27.9***
[N=122]	[-36.5:-1.0]	[-44.0:-12.0]

Reported median difference row and column values

Comment: [Confidence interval at the appropriate significance level (def. 90%)]

***/**/* denote rejection of H_0 at 1/5/10% significance levels

†/‡ denote rejection of H_0 at 5/10% significance levels before Holm-Bonferroni correction

Table 4.8: Combinatorial curse

Auction	Affected wins	Profit	Potential profit	Potential gain
First-price	35 out of 93	38.3	46.2	7.9
Vickrey	35 out of 97	61.9	76.4	14.5
VNR	33 out of 107	51.3	66.2	14.9
Ref. Rule	37 out of 97	58.2	76.0	17.8

Comment: Average profit conditional on winning

in Vickrey and MRCS auctions, most of the local bidders' bids are equal to their value, hence shading is zero. This creates a lot of ties and makes the Mann-Whitney test very sensitive to the tail of the shading distribution. For example, in the Components treatment of the Vickrey auction, local bidders shaded just 17 out of 140 times on item x (mean shading=0.85). In the Bundle treatment, these bidders shaded just 5 out of 150 (mean shading=0.086). Yet the Mann-Whitney test rejects the null hypothesis (adj. p -value=0.007; unadj. p -value=0.205) that these distributions are identical.³⁰

Therefore, we only compare the raw bid data across Components/Bundle and Exposed/Bundle treatments and auctions for the local bidders using pairwise Mann-Whitney tests. This test should be isomorphic to testing shading data as long as the realised value distributions are the same for all local bidders. However, there is no systematic pattern of treatment effects on bidders. In fact, for every auction, we cannot reject the null hypothesis of equality for any treatment (see Appendix 4.7.2.3). While local bidders bid differently across auctions, we have no conclusive evidence that they bid differently across treatments. Therefore, in line with Bykowsky et al. (2000), we find no evidence of threshold bidding behaviour even though in our setting local bidders could not communicate. It is likely that the presence of threshold effects in first-price auctions conducted by Chernomaz and Levin (2012) is due to the ex post symmetry of local bidders. Knowing the valuation of the other local bidder makes free-riding more appealing.

E. Local bidders bid their value in the Vickrey auction. Overall, there is a lot of truthful bidding in our data. However, in the Vickrey auction, truthful bidding is so

³⁰One option would be to use Mood's non-parametric test for equality of medians, which allows for different tie-breaking patterns. However, this test has low power and its conclusions are economically uninteresting in our context. Another option is the Ansari-Bradley test for the difference in dispersions of two distributions with the same median (Ansari and Bradley, 1960). However, this test is too sensitive because the variances are affected by a small number of values. Finally, instead of shading=value-bid, we could look at bid ratio= $\frac{\text{bid}}{\text{value}}$, however, this measure can be sensitive at low realisations of the value distribution.

Table 4.9: Frequency of truthful bidding in the Vickrey auction

	Local bidders' bids					
	On item x			On item y		
	Truthful	Non-truthful	$\frac{\text{Truthful}}{\text{Total}}$	Truthful	Non-truthful	$\frac{\text{Truthful}}{\text{Total}}$
Exposed	133	7	0.95	131	9	0.94
Components	123	17	0.88	118	22	0.84
Bundle	135	5	0.96	133	7	0.95
Combinatorial	133	7	0.95	133	7	0.95

pervasive among local bidders that the small variance in the data does not allow us to use any statistical tests or models to test for equality of bids and values. The Vuong test does not reject the equality of a zero-inflated and non-inflated negative binomial models in any treatment.³¹ Hence, we simply report the proportion of truthful bids in the Vickrey auction, which we hope will convince the reader that the prediction of the Vickrey auction are borne out in our data. Table 4.9 reports the frequency of truthful and non-truthful bidding by local bidders in the Vickrey auction.

F. Local bidders shade less in the Reference Rule than in the Vickrey-Nearest Rule. Pairwise Mann-Whitney tests on the raw bid data cannot reject the null hypothesis that local bidders bid identically in the Reference Rule and Vickrey-Nearest Rule auctions across all treatments (see Appendix 4.7.2.4). This happens despite the fact that bidders' payments are sufficiently asymmetric to be able to identify differences if there were any in the Bundle and Combinatorial treatments. Therefore, we do not find any evidence that bidders have responded to the lower sum of marginal incentives to shade in the Reference Rule compared to the Vickrey-Nearest Rule auction.

In addition, we compared the response to incentives by Bidders 1 and 2 in the Reference Rule auction. Whenever payments are equal to the reference point, Bidder 1 pays three times more for his item. This should give him a stronger incentive to shade

³¹Once again, this could either be caused by small non-zero samples or by the low power of the Vuong test in finite samples.

compare to the Bidder 2. Pairwise Mann-Whitney tests on raw bid data provide no evidence of any difference, however, when we run them on the sensitive shading data, the null hypothesis of equality is rejected in all but the Combinatorial treatment. We also relegate the tables to Appendix 4.7.2.4.

4.5.3 Discussion

Our results offer interesting perspectives on the theoretical auctions literature. Some of the theoretical predictions, such as local bidders' truthful bidding in the Vickrey auction, get excellent support in the data. However, there are many challenges. We observe unrivaled dominance of the first-price auction in terms of revenue without much loss in efficiency. It is a well-established fact in the auctions literature that subjects overbid in first-price auctions relative to equilibrium behaviour (Kagel and Levin, 2011). Marszalec (2011) found that subjects also overbid in MRCS auctions and we can reasonably conjecture that they also overbid in MRCS auctions in our settings. However, the marginal effect of overbidding on revenue in the first-price auction is greater than in MRCS auctions because the payment depends entirely on the bid. This effect drives the robustness of revenue-dominance of the first-price auction. While the reasons for overbidding are numerous – risk aversion (Cox et al., 1988), impulse and social comparison (Ockenfels and Selten, 2005, Neugebauer and Selten, 2006)³², and hierarchical non-equilibrium strategic behaviour (Crawford and Iriberry, 2007)³³ – they are not the focus of this chapter.

In MRCS auctions, local bidders do not respond to incentives to shade and bid similarly to Vickrey auctions in Bundle and Combinatorial treatments. We cannot reject the null hypothesis that local bidders bid differently in the Vickrey auction and

³²Although in our setting, subjects receive plenty of feedback (see Figure 4.5b), therefore we cannot attribute overbidding to lack to feedback.

³³Truthful bidding in complex MRCS auctions is reasonable starting point for a hierarchical behavioural model of strategic reasoning (i.e. truthful $L0$). Hence, many of the subjects who shade in the MRCS auction might be best-responding to this strategy making them truthful $L1$ -type bidders (Crawford and Iriberry, 2007).

MRCS auctions (see Appendix 4.7.2.4). This is, of course, the theoretical prediction in Exposed and Components treatments. This is good news for the MRCS auctions because, in general, they should revenue-dominate the Vickrey auction without much loss in efficiency in package treatments. However, we have not found any noticeable differences between the Vickrey-Nearest Rule and the Reference Rule auctions. In the absence of substantial differences between these two auctions, Vickrey-Nearest Rule may be preferred by practitioners because the relative payments in the MRC are fully endogenously determined by the bids.³⁴ On the other hand, a Reference Rule auction may be justified for equity considerations.

It is not clear whether the efficiency of the first-price auction or the absence of differences between the MRCS auctions will persist when there are more items in the auction. Many real-world auctions, such as spectrum auctions, have tens of items for sale. In fact, bidder behaviour in very large-value, one-off auctions may differ substantially from our results. However, the conclusions of this chapter cannot be dismissed entirely.

Our subjects appear to be adversely affected by the complexity of auction design. In the Combinatorial setting, especially in the more computationally-intensive MRCS auction, subjects appear to take mental shortcuts and report their true bids thereby reducing their total profits. They do not take the effect of their bids fully into account. It is not clear whether this result will transfer to dynamic settings.³⁵

³⁴The auctioneer does not need to justify his use of the reference point.

³⁵For the two-item three-bidder setting discussed in this chapter, competitive equilibrium always exists because bidder valuations satisfy the gross substitutes and complements structure (Sun and Yang, 2006). The dynamic, double-track process described in Sun and Yang (2009) will always find an efficient allocation and prices that support it. Sun and Yang (2008) show that there exists a strategy-proof mechanism based on the double-track process that implements the efficient allocation. We discuss this auction further in Chapter 3 of this thesis. Goeree and Lien (2009) provide a Bayesian equilibrium analysis of the dynamic, non-package simultaneous multiple round (SMR) auction originally used by the FCC in a setting with multiple local and global bidders.

4.6 Conclusions and extensions

This chapter looked at the effect of exposure and package bidding in four auctions with complements. We considered a simple two-item setting with three bidders. We found that the first-price auction dominates the other auctions in revenue without any loss in efficiency. Global bidders are affected by exposure only in the first-price auction. We found no threshold effects or substantial differences between the outcomes or bidding in Reference Rule and Vickrey Nearest Rule auctions. This chapter complements previous experiments by Marszalec (2011) as well as other experiments we have done together. The main policy-relevant conclusion is that, despite some theoretical concerns, first-price sealed-bid auctions revenue-dominate (a) sealed-bid MRCS package auctions in settings with strong complementarities and (b) second-price sealed-bid auctions in non-package settings with exposure *without any efficiency loss*.

One clear direction for further research in auctions with complements is theoretical. While it would be useful to find equilibrium solutions to the auction settings presented here, they are unlikely to offer much practical guidance since experimental subjects do not conform to equilibrium behaviour in even simpler settings. Instead it would be interesting to incorporate the experimental findings of our research into a richer theory of strategic behaviour in complex environments. Some advances in this direction have already been made (Crawford et al., 2009, Jehiel, 2011), but much work remains.

There are also obvious areas for further experimental work. It is not clear how competition between bidders and increasing the number of items on sale will affect revenue and efficiency. Rigorously testing these environments in the experimental lab would allow us to see whether the first-price auction is more suited in supplemental sealed-bid phases of the spectrum auctions than MRCS auctions. Yet our experience suggests that even a three-item combinatorial setting may be too difficult for subjects

to test the Vickrey or MRCS auctions in a lab environment. Finally, it would also be interesting to test experimentally the performance other core-selecting auctions such as the Nearest-Bid, Proportional Rule, and Proxy Rule auctions against their theoretical benchmarks (Ausubel and Baranov, 2010).

4.7 Appendix

4.7.1 Subject characteristics and test performance

Table 4.10: Subject characteristics

Characteristic	Proportion
Gender	
<i>Male</i>	0.60
<i>Female</i>	0.40
English as first language	0.59
Status	
<i>Undergraduate</i>	0.53
<i>Postgraduate</i>	0.45
<i>Other</i>	0.01
Degree subject	
<i>STEM/Medicine</i>	0.52
<i>Social Sciences</i>	0.40
<i>Humanities</i>	0.09

Proportions may not add up to 1 due to rounding.

STEM = Science, Technology, Engineering, and Mathematics.

Table 4.11: Test results

Session	Auction Type	Total test passes	Total test fails	Quit
1	First-price	24	1	0
2	Ref. Rule	24	1	0
3	VNR	21	2	1
4	Vickrey	24	1	0
5	First-price	18	3	0
6	Ref. Rule	18	4	0
7	VNR	12	13	1
8	VNR	15	10	1
9	Vickrey	18	10	1

4.7.1.1 Examples of bidding and feedback screens

Figure 4.5: Examples of z-Tree software

(a) An example of a bidding screen

Period	1 of 10	Remaining time [sec]: 8
Your Bidder Type is: 3 You can bid on the bundle X & Y, and also on X and Y separately.		
Your value on item X alone is: 57.0 Your value on item Y alone is: 93.0 Your value on the bundle X & Y together is: 200.0 You can bid up to - and including - your value.		
Your bid on X is: <input type="text"/> Your bid on Y is: <input type="text"/> Your bid on the bundle (X & Y) is: <input type="text"/>		
<input type="button" value="Continue"/>		

(b) An example of a feedback screen

Period	1 of 10	Remaining time [sec]: 5
Your type is: 3 Your bid on X alone was: 40.0 Your bid on Y alone was: 60.0 Your bid on the bundle X & Y was: 170.0 Your bid WAS part of the winning allocation ! You won both items X & Y through a bundle bid.		
The highest bid on X was: 73.0 The highest bid on Y was: 60.0 The highest bid on the bundle (X & Y) was: 170.0 The sum of the winning bids was: 170.0 The underbid was: 133.0		
Your VALUE for the items X & Y was: 200.0 Your PAYMENT for this period was: 133.0 Your PROFIT from this round is: 67.0		
<input type="button" value="Continue"/>		

4.7.2 Further tables

4.7.2.1 Effect of auction design on bidder surplus and profits

Table 4.12: Median-difference test results — Surplus

Exposed	Vickrey [N=140]	First-price [N=140]	VNR [N=160]
First-price [N=140]	−56.5*** [−72.0:−41.5]	—	—
VNR [N=160]	0.0 [−10.5:11.0]	56.0*** [43.0:71.0]	—
Ref. Rule [N=140]	0.0 [−12.0:10.5]	57.0*** [43.5:71.0]	−0.5 [−11.0:10.0]
Components	Vickrey [N=140]	First-price [N=140]	VNR [N=160]
First-price [N=140]	−41.0*** [−52.5:−30.0]	—	—
VNR [N=160]	1.0 [−8.0:10.0]	40.0*** [29.0:52.0]	—
Ref. Rule [N=140]	4.0 [−5.0:14.0]	46.0*** [34.5:57.0]	4.0 [−6.0:13.0]
Bundle	Vickrey [N=140]	First-price [N=140]	VNR [N=160]
First-price [N=140]	−38.0*** [−49.5:−26.0]	—	—
VNR [N=160]	−7.0 [−16.0:1.5]	29.5*** [16.5:43.5]	—
Ref. Rule [N=140]	−4.0 [−13.0:5.0]	32.0*** [20.5:44.0]	4.0 [−4.5:13.0]
Combinatorial	Vickrey [N=140]	First-price [N=140]	VNR [N=160]
First-price [N=140]	−20.5*** [−34.5:−8.0]	—	—
VNR [N=160]	−7.5 [−16.0:0.5]	12.0*** [2.5:22.5]	—
Ref. Rule [N=140]	−6.0 [−15.0:3.0]	14.5*** [3.0:26.0]	2.0 [−6.5:10.0]
Reported median difference between row and column values			
<i>Comment:</i>	[Confidence interval at the appropriate significance level (default 90%)]		
	***/**/* denote rejection of H_0 at the 1/5/10% significance levels		
	No tests are affected by Holm-Bonferroni correction		

Table 4.13: Median-difference test results — Global bidders' profit

Exposed	Vickrey [N=140]	First-price [N=140]	VNR [N=160]
First-price [N=140]	−49.0*** [−65.0:−32.0]	—	—
VNR [N=160]	0.0 [−9.0:9.0]	51.0*** [35.0:66.0]	—
Ref. Rule [N=140]	0.0 [−12.5:9.0]	50.0*** [30.1:66.0]	0.0 [−12.0:9.0]
Components	Vickrey [N=140]	First-price [N=140]	VNR [N=160]
First-price [N=140]	−24.0*** [−39.0:−8.0]	—	—
VNR [N=160]	0.0 [−8.0:8.0]	22.0*** [6.5:39.0]	—
Ref. Rule [N=140]	0.0 [−6.0:10.0]	27.0*** [9.0:44.250]	0.0 [−5.0:10.0]
Bundle	Vickrey [N=140]	First-price [N=140]	VNR [N=160]
First-price [N=140]	−19.5*** [−35.5:−1.5]	—	—
VNR [N=160]	0.0 [−3.5:5.5]	18.0*** [1.5:34.5]	—
Ref. Rule [N=140]	1.0 [0.0:14.0]	26.0*** [10.0:41.5]	1.0 [0.0:14.0]
Combinatorial	Vickrey [N=140]	First-price [N=140]	VNR [N=160]
First-price [N=140]	−8.0 [−17.0:0.0]	—	—
VNR [N=160]	−5.5 [−14.0:0.0]	0.0 [−5.0:4.0]	—
Ref. Rule [N=140]	0.0 [−10.5:5.625]	6.5 [0.0:14.0]	4.5 [0.0:12.5]
Reported median difference between row and column values			
<i>Comment:</i>	[Confidence interval at the appropriate significance level (default 90%)]		
	***/**/* denote rejection of H_0 at the 1/5/10% significance levels		
	No tests are affected by Holm-Bonferroni correction		

4.7.2.2 Profits and revenues in the Combinatorial treatment

Table 4.14: Median-difference test results — Global bidders' profit

	Components [N=140]	Bundle [N=140]
Vickrey		
Combinatorial [N=140]	-2.0 [-13.0:2.0]	-3.5 [-16.0:0.0]
First-price		
Combinatorial [N=140]	6.0** [1.0:11.0]	0.0 [-3.0:4.5]
VNR		
Combinatorial [N=160]	-10.5** [-19.0:0.0]	-13.0*** [-25.5:-1.5]
Ref. Rule		
Combinatorial [N=140]	-7.0 [-18.5:0.0]	-15.5*** [-29.0:-2.0]
<p>Reported median difference between row and column values</p> <p><i>Comment:</i> [Confidence interval at the appropriate significance level (default 90%)</p> <p>***/**/* denote rejection of H_0 at the 1/5/10% significance levels</p> <p>No tests are affected by Holm-Bonferroni correction</p>		

Table 4.15: Median-difference test results — Revenue

	Components [N=140]	Bundle [N=140]
Vickrey		
Combinatorial [N=140]	20.0*** [7.0:33.0]	18.0*** [4.0:32.0]
First-price		
Combinatorial [N=140]	0.0 [-5.0:6.0]	3.0 [-3.0:9.0]
VNR		
Combinatorial [N=160]	27.0*** [15.0:40.0]	17.5*** [5.5:30.0]
Ref. Rule		
Combinatorial [N=140]	21.5*** [9.0:35.0]	9.0* [1.0:17.0]
<p>Reported median difference between row and column values</p> <p><i>Comment:</i> [Confidence interval at the appropriate significance level (default 90%)</p> <p>***/**/* denote rejection of H_0 at the 1/5/10% significance levels</p> <p>No tests are affected by Holm-Bonferroni correction</p>		

4.7.2.3 Threshold bidding hypothesis

Table 4.16: Mann-Whitney Test — Bid on item X — Local bidder 1

	Exposed	Components
Vickrey		
Bundle [N=140]	9.0 [52.5:43.5] 0.130	9.0 [†] [52.5:43.5] 0.094
First-price		
Bundle [N=140]	13.8 [40.0:26.3] 0.176	0.0 [40.0:40.0] 0.856
VNR		
Bundle [N=160]	-4.0 [46.0:50.0] 0.392	-4.0 [46.0:50.0] 0.571
Ref. Rule		
Bundle [N=140]	-2.5 [45.0:47.5] 0.884	-4.5 [45.0:49.5] 0.879
<p>Reported difference between row median and column median</p> <p><i>Comment:</i> [median(row):median(column)] — p-value calculated using z-score</p> <p>***/**/* denote rejection of H_0 at the 1/5/10% significance levels</p> <p>†/‡ denote rejection of H_0 at 5/10% significance levels before Holm-Bonferroni correction</p>		

Table 4.17: Mann-Whitney Test — Bid on item Y — Local bidder 2

	Exposed	Components
Vickrey		
Bundle [N=140]	3.0 [49.0:46.0] 0.541	-2.0 [49.0:51.0] 0.354
First-price		
Bundle [N=140]	1.0 [41.0:40.0] 0.841	0.5 [41.0:40.5] 0.545
VNR		
Bundle [N=160]	-1.0 [52.0:53.0] 0.691	8.0 [52.0:44.0] 0.294
Ref. Rule		
Bundle [N=140]	1.5 [49.5:48.0] 0.979	2.0 [49.5:47.5] 0.730
<p>Reported difference between row median and column median</p> <p><i>Comment:</i> [median(row):median(column)] — p-value calculated using z-score</p> <p>***/**/* denote rejection of H_0 at the 1/5/10% significance levels</p> <p>†/‡ denote rejection of H_0 at 5/10% significance levels before Holm-Bonferroni correction</p>		

4.7.2.4 Vickrey-Nearest Rule vs. Reference Rule

Table 4.18: Mann-Whitney Test — Bid on item X — Local bidder 1

Exposed	Vickrey [N=140]	First-price [N=140]	VNR [N=160]	Ref. Rule [N=160]
Vickrey [N=140]	—	—	—	—
First-price [N=140]	-17.3*** [26:44] 0.002	—	—	—
VNR [N=160]	6.5 [50:44] 0.297	23.8*** [50:26] 0.000	—	—
Ref. Rule [N=160]	4.0 [48:44] 0.654	21.3*** [48:26] 0.000	-2.5 [48:50] 0.544	—
Components	Vickrey [N=140]	First-price [N=140]	VNR [N=160]	Ref. Rule [N=160]
Vickrey [N=140]	—	—	—	—
First-price [N=140]	-3.5* [40:44] 0.064	—	—	—
VNR [N=160]	6.5 [50:44] 0.392	10.0*** [50:40] 0.004	—	—
Ref. Rule [N=160]	6.0 [50:44] 0.433	9.5*** [50:40] 0.006	-0.5 [50:50] 0.991	—
Bundle	Vickrey [N=140]	First-price [N=140]	VNR [N=160]	Ref. Rule [N=160]
Vickrey [N=140]	—	—	—	—
First-price [N=140]	-12.5*** [40:53] 0.000	—	—	—
VNR [N=160]	-6.5 [46:53] 0.150	6.0** [46:40] 0.023	—	—
Ref. Rule [N=160]	-7.5 [45:53] 0.268	5.0*** [45:40] 0.006	-1.0 [45:46] 0.691	—
Combinatorial	Vickrey [N=140]	First-price [N=140]	VNR [N=160]	Ref. Rule [N=160]
Vickrey [N=140]	—	—	—	—
First-price [N=140]	-20.0*** [39:59] 0.000	—	—	—
VNR [N=160]	-3.5 [56:59] 0.519	16.5*** [56:39] 0.000	—	—
Ref. Rule [N=160]	-12.0** [47:59] 0.046	8.0** [47:39] 0.010	-8.5 [47:56] 0.183	—
Reported difference between row median and column median				
<i>Comment:</i>	[median(row):median(column)] — <i>p</i> -value calculated using z-score			
	***/**/* denote rejection of H_0 at the 1/5/10% significance levels			
	uncorrected by the Holm-Bonferroni method			

Table 4.19: Mann-Whitney Test — Bid on item Y — Local bidder 2

Exposed	Vickrey [N=140]	First-price [N=140]	VNR [N=160]	Ref. Rule [N=160]
Vickrey [N=140]	—	—	—	—
First-price [N=140]	−6.0** [40:46] 0.035	—	—	—
VNR [N=160]	7.0 [53:46] 0.310	13.0*** [53:40] 0.001	—	—
Ref. Rule [N=160]	2.0 [48:46] 0.640	8.0*** [48:40] 0.006	−5.0 [48:53] 0.640	—
Components	Vickrey [N=140]	First-price [N=140]	VNR [N=160]	Ref. Rule [N=160]
Vickrey [N=140]	—	—	—	—
First-price [N=140]	−10.5*** [41:51] 0.000	—	—	—
VNR [N=160]	−7.0** [44:51] 0.045	3.5 [44:41] 0.144	—	—
Ref. Rule [N=160]	−3.5 [48:51] 0.129	7.0** [48:41] 0.032	3.5 [48:44] 0.615	—
Bundle	Vickrey [N=140]	First-price [N=140]	VNR [N=160]	Ref. Rule [N=160]
Vickrey [N=140]	—	—	—	—
First-price [N=140]	−8.0*** [41:49] 0.002	—	—	—
VNR [N=160]	3.0 [52:49] 0.969	11.0*** [52:41] 0.001	—	—
Ref. Rule [N=160]	0.5 [50:49] 0.840	8.5*** [50:41] 0.005	−2.5 [50:52] 0.907	—
Combinatorial	Vickrey [N=140]	First-price [N=140]	VNR [N=160]	Ref. Rule [N=160]
Vickrey [N=140]	—	—	—	—
First-price [N=140]	−1.0 [42:43] 0.114	—	—	—
VNR [N=160]	8.3 [51:43] 0.238	9.3*** [51:42] 0.001	—	—
Ref. Rule [N=160]	−2.0 [41:43] 0.573	−1.0 [41:42] 0.215	−10.3 [41:51] 0.103	—
Reported difference between row median and column median				
<i>Comment:</i>	[median(row):median(column)] — <i>p</i> -value calculated using z-score			
	***/**/* denote rejection of H_0 at the 1/5/10% significance levels			
	uncorrected by the Holm-Bonferroni method			

Table 4.20: Differences in local bidders' bids in the Reference Rule

	Median bid of bidder		Difference	<i>p</i> -value
	1 on item X	2 on item Y		
Exposed	47.5	48	-0.5	0.624
Components	49.5	47.5	2	0.841
Bundle	45	49.5	-4.5	0.823
Combinatorial	47.5	40.5	7.5	0.430

Comment: *p*-value for a two-sided Mann-Whitney test

4.7.3 Instructions

The following pages present instructions and the test given to participants during the Vickrey auction experimental session. The solution to the understanding test is at the end. All instructions were read aloud by the supervisor. Instructions for other auctions are available upon request.

Instructions for the Auctions Experiment

You are about to participate in an experiment on auctions. The university and research foundations have provided funds for the experiment. If you follow the instructions, and pass the Understanding Tests (in the first stage of the experiment), you will be allowed to continue.

Depending on your decisions, you may then earn a considerable amount of money, from £8 to £68 (on average, around £40). The duration of this experiment will be 2 to 3 hours.

Your additional earnings will be determined by your decisions and the decisions of other participants in the experiment.

All the money that you earn during the experiment is yours to keep, and will be paid to you in private, in cash, after today's session. The experiment will consist of **4 sections**, in each of which you will bid for goods against other participants in the room in a series of auctions. When all subjects have completed the **10 auction rounds** in a section, you will be able to go on to the next section.

In the session today your and the others' bids will determine whether you win an item, and how much the winner(s) pays for the object(s).

In each auction round of each section, you will be grouped with two other participants, chosen randomly from the room. None of you will ever know with whom you have been grouped, and the groups will change randomly from round to round. The rules for calculating the winners' payments from the bids of all the bidders in their group will not vary from section to section.

How the auction works is explained at the beginning of this session and you will be permitted to ask clarifying questions. You will then complete an Understanding Test, and two practice rounds to make sure that you have understood the rules of the auction. The results of the test rounds will not influence your payment, but if you fail the Understanding Test, you will not be allowed to participate.

After the two test rounds, you will participate in 10 rounds of the auction run under the same conditions. In each auction, the computer will show you your value for an item or items, and then request you to submit a bid. The value that you are assigned will be drawn independently for each person in each auction round, and each section. So your value, and the values of the other bidders, will change throughout the experiment.

At the end of each section, two rounds from that section will be selected at random, and you will be paid for that section according to your profits in those two rounds. If your profit in those two rounds is negative, it will be capped at zero. Note that to maximise your earnings under this compensation scheme, you need to try to maximise your profit in each individual round - it is *not* sufficient that you 'do well' in a few rounds only, since those rounds may not get chosen by the computer.

The exchange rate of points to Pounds is: 5 points to £1. These profits will be paid out to you in cash, in private, at the end of the experiment.

In each auction there will be two items to bid on: item X and item Y. These are virtual items. You will be randomly assigned to one of three 'Types', which will determine your value for these two items. Your Type may change from one auction round to the next. This will be made clear at the beginning of each section.

The **Types** are as follows:

In all sections of the auction, Type 1 has a positive value for item X and a zero value for item Y. Her value for item for X is a random number from 0 to 100, with each number equally likely.

In all sections of the auction, Type 2 has a positive value for item Y and a zero value for item X. Her value for item for Y is a random number from 0 to 100, with each number equally likely.

In all sections of the auction, Type 3 has a positive value for the bundle of items X & Y together. Her value for the individual items X and Y and for the bundle of items X & Y will vary from section to section and will be made clear at the beginning of each section.

Bidding rules

In all sections of the auction, Type 1 will be allowed to bid for item X only, and only up to (and including) her value for item X.

In all sections of the auction, Type 2 will be allowed to bid for item Y only, and only up to (and including) her value for item Y.

Type 3 will be allowed to bid for the bundle of items X & Y, and/or on item X, and/or on item Y depending on the section of the auction. This will be made clear at the beginning of each section.

Hence, Type 3 may be able to win the bundle of items X & Y either by submitting winning bids for item X and item Y separately or submitting a winning bid for the bundle of items X & Y.

In all sections of the auction, Type 3 will be allowed to bid on any individual item or on the bundle up to (and including) her value of the bundle X & Y.

Practicalities

Please keep this instruction sheet with you throughout the experiment.

At the end of each section, please note down your 'payment reference' number in the space below, so that we can verify that you are paid correctly:

Payment Reference Number in Section A: _____

Payment Reference Number in Section B: _____

Payment Reference Number in Section C: _____

Payment Reference Number in Section D: _____

Terminal/Cubicle Number (& letter, if applicable): _____

The Rules of the Auction

In this auction the winners of the items will be determined from the participants' bids to make the sum total of the winning bids as high as possible. Thus if the highest sum of individual bids on item X and on item Y is greater than any bids on the bundle of items X & Y, then bidder(s), who made those bids on the individual items, win them. If the highest bid on the bundle of items X & Y exceeds the highest sum of individual bids on item X and item Y, then the Type 3 bidder wins both items X & Y (because only the Type 3 bidder can bid on the bundle). Hence, in some sections of the auction Type 3 may be able to win the bundle of items X & Y either by submitting winning bids for item X and item Y separately or submitting a winning bid for the bundle of items X & Y. If the sum of the highest individual bids on items X and Y equals the highest bid on the bundle of items X & Y, then the tie will be broken by a 'fair coin toss' (simulated by the computer).

If you do not win an item in a round of the auction, your profit for that round will be zero. If you do win an item, your **profit** will be your **value** for that item(s), minus your **payment(s)** for the item(s) that you won. If your payment exceeds your value, then your profit will be negative.

In this auction your own payment for the item(s) – if you win – will only depend on bids of other bidders.

If you **do not win** any items you will **pay zero** and your **profit** will be **zero**.

Winning either item X or item Y

In order to calculate your payment when you win:

- Determine the **other winning bid**.
- Using current bids, find what bidders would have won the auction if *your winning bid alone were absent* from the auction. Note that the auctioneer does not need to sell both items in this case.
- Calculate the highest sum of winning bids from the remaining bids if *your winning bid alone were absent* from the auction. Call this **absent revenue (AR)**.
- Your **payment for the item** is the **absent revenue** minus the **other winning bid**.
- Your **profit** is the difference between your **value** and your **payment**

Winning both item X and item Y

In some sections of the auction, Type 3 may be able to win the bundle of items X & Y either by submitting winning bids for item X and item Y separately or submitting a winning bid for the bundle of items X & Y. To calculate the **total payment if you win both items X and Y**

- Using current bids, find what bidders would have won the auction if *your winning bid(s) were absent* from the auction. Note that the auctioneer does not need to sell both items in this case.
- Calculate the highest sum of winning bids from the remaining bids if *your winning bid(s) were absent* from the auction. Call this the **absent revenue (AR)**.
- Your **payment for the items** is the **absent revenue**.
- Your **profit** is the difference between your **value** and your **payment**.

Section A

In this section, the Type 3 bidder can only bid on item X and/or item Y individually. Her value for item X is a random number between 0 and 100, with each number equally likely, and her value for item Y is a random number between 0 and 100, with each number equally likely.

Type 3 bidder's value for the bundle of items X & Y is the sum of her values on the individual items X and Y multiplied by 1.5, but at most 200. Succinctly, it is: $\min([\text{value for X} + \text{value for Y}] \times 1.5, 200)$.

Example 1

Bidder type	Value...	...on item(s)	Bid	Winning bid?	Absent Revenue	Payment	Total Profit
1	82	X	56	Yes	92	50	32
2	60	Y	42	Yes	90	34	26
3	35	X	50	No	-	0	0
3	65	Y	34	No	-	0	0
3	150	X & Y	N/A	-	-	0	0

Determining the winner(s)

Type 1 bidder wins item X and Type 2 bidder wins item Y because the sum of their bids ($98=56+42$) is greater than the sum of any other bids on the individual items.

Calculating payments

For Type 1 bidder: **Other winning bid** is 42.

Absent revenue is $92 = 42 + 50$ because bidder Type 2 would have won item Y and bidder Type 3 would have won item X if Type 1's winning bid were absent from the auction.

Type 1's **payment** is $92 - 42 = 50$. The **profit** of Type 1 bidder is $82-50=32$.

For Type 2 bidder: **Other winning bid** is 56.

Absent revenue is $90 = 56 + 34$ because bidder Type 1 would have won item X and bidder Type 3 would have won item Y if Type 2's winning bid were absent from the auction.

Type 2's **payment** is $90 - 56 = 34$. The **profit** of Type 2 bidder is $60-34=26$.

For Type 3 bidder: She does not win, and therefore pays 0, and gets profit 0.

Example 2

Bidder type	Value...	...on item(s)	Bid	Winning bid?	Absent Revenue	Payment	Total Profit
1	31	X	30	No	-	0	0
2	50	Y	46	Yes	90	40	10
3	10	X	50	Yes	76	30	-20
3	50	Y	40	No	-	0	0
3	90	X & Y	N/A	-	-	0	0

Determining the winner(s)

Type 2 wins item Y and Type 3 wins item X because the sum of their bids ($96 = 46 + 50$) is greater than the sum of any other bids.

Calculating payments

For Type 1 bidder: She does not win, and therefore pays 0, and gets profit 0.

For Type 2 bidder: **Other winning bid** is 50.

Absent revenue is $90 = 50 + 40$ because bidder Type 3 would have won both items if Type 2's winning bid were absent from the auction.

Type 2's **payment** is: $90 - 50 = 40$. The **profit** of Type 2 bidder is $50 - 40 = 10$.

For Type 3 bidder: **Other winning bid** is 46.

Absent revenue is $76 = 46 + 30$ because bidder Type 1 would have won item X and bidder Type 2 would have won item Y if Type 3's winning bid were absent from the auction.

Type 3's **payment** is: $76 - 36 = 30$. The **profit** of Type 3 bidder is $10 - 30 = -20$.

Example 3

Bidder type	Value...	...on item(s)	Bid	Winning bid?	Absent Revenue	Payment	Total Profit
1	24	X	13	No	-	0	0
2	66	Y	47	No	-	0	0
3	50	X	95	Yes	↓	↓	↓
3	70	Y	95	Yes	↓	↓	↓
3	180	X & Y	N/A	-	60	60	120

Determining the winner(s)

Type 3 wins both item X and item Y because the sum of her bids ($190 = 95 + 95$) is greater than the sum of any other bids.

Calculating payments

For Type 1 bidder: She does not win, and therefore pays 0, and gets profit 0.

For Type 2 bidder: She does not win, and therefore pays 0, and gets profit 0.

For Type 3 bidder: **Absent revenue** is $60 = 13 + 47$ because Type 1 would have won item X and Type 2 would have won item Y if Type 3's winning bids on item X and on item Y were absent from the auction.

Type 3's **payment** is the **absent revenue** of 60. The **profit** of Type 3 bidder is $180 - 60 = 120$.

Scrap paper

Understanding Test

Participant Number: _____

The four tables below are to test your understanding of the Rules of the Auction. The table provides you with the values for each bidder type, and the actual bids that each type has submitted. To confirm that you understand how the rules of this auction work, you need to fill in the following three details for each bidder type:

- Does this type 'win' an object?
- What is this type's payment?
- What is this type's profit?

In completing this test, you ARE permitted to refer to the instruction sheet and/or use the calculator.

Table 1

Bidder type	Value...	...on item(s)	Bid	Winning bid?	Absent Revenue	Payment	Total Profit
1	36	X	26				
2	70	Y	66				
3	35	X	14				
3	65	Y	28				
3	150	X & Y	N/A				

Table 2

Bidder type	Value...	...on item(s)	Bid	Winning bid?	Absent Revenue	Payment	Total Profit
1	36	X	31				
2	70	Y	33				
3	35	X	34				
3	65	Y	90				
3	150	X & Y	N/A				

Table 3

Bidder type	Value...	...on item(s)	Bid	Winning bid?	Absent Revenue	Payment	Total Profit
1	32	X	25				
2	70	Y	16				
3	35	X	20				
3	65	Y	45				
3	150	X & Y	N/A				

Table 4

Bidder type	Value...	...on item(s)	Bid	Winning bid?	Absent Revenue	Payment	Total Profit
1	87	X	85				
2	90	Y	47				
3	10	X	25				
3	50	Y	80				
3	90	X & Y	N/A				

Section B

In this section, the Type 3 bidder can bid on item X and item Y individually and on the bundle of items X & Y. Her value for item X is a random number between 0 and 100, with each number equally likely, and her value for item Y is a random number between 0 and 100, with each number equally likely.

Type 3 bidder's value for the bundle of items X & Y is the sum of her values on the individual items X and Y multiplied by 1.5, but at most 200. Succinctly, it is: $\min([\text{value for X} + \text{value for Y}] \times 1.5, 200)$.

Example 4

Bidder type	Value...	...on item(s)	Bid	Winning bid?	Absent Revenue	Payment	Total Profit
1	82	X	56	No	-	0	0
2	54	Y	40	No	-	0	0
3	35	X	100	Yes	↓	↓	↓
3	65	Y	100	Yes	↓	↓	↓
3	150	X & Y	120	No	120	120	30

Determining the winner(s)

Type 3 bidder wins both items because the sum of her individual bids ($200 = 100 + 100$) is greater than the sum of any other bids on the individual items.

Calculating payments

For Type 1 bidder: She does not win, and therefore pays 0, and gets profit 0.

For Type 2 bidder: She does not win, and therefore pays 0, and gets profit 0.

For Type 3 bidder: **Absent revenue** is 120 because Type 3 would have won both items with the bid on the bundle of items X & Y if the winning bids of Type 3 were absent from the auction.

Type 3's **payment** is the **absent revenue** of 120. The **profit** of Type 3 is $150 - 120 = 30$.

Example 5

Bidder type	Value...	...on item(s)	Bid	Winning bid?	Absent Revenue	Payment	Total Profit
1	82	X	37	Yes	66	32	50
2	54	Y	34	Yes	40	3	51
3	10	X	32	No	-	0	0
3	20	Y	2	No	-	0	0
3	45	X & Y	40	No	-	0	0

Determining the winner(s)

Type 1 wins item X and Type 2 wins item Y because the sum of their bids ($71 = 37 + 34$) is greater than the sum of any other bids.

Calculating payments

For Type 1 bidder: **Other winning bid** is 34.

Absent revenue is $66 = 34 + 32$ because bidder Type 2 would have won item Y and bidder Type 3 would have won item X if Type 1's winning bid were absent from the auction.

Type 1's **payment** is $66 - 34 = 32$. The **profit** of Type 1 bidder is $82 - 32 = 50$.

For Type 2 bidder: **Other winning bid** is 37.

Absent revenue is 40 because bidder Type 3 would have won both items with her bid on the bundle of items X & Y if Type 2's winning bid were absent from the auction.

Type 2's **payment** is $40 - 37 = 3$. The **profit** of Type 2 bidder is $54 - 3 = 51$.

For Type 3 bidder: She does not win, and therefore pays 0, and gets profit 0.

Example 6

Bidder type	Value...	...on item(s)	Bid	Winning bid?	Absent Revenue	Payment	Total Profit
1	82	X	82	No	-	0	0
2	54	Y	20	Yes	128	18	36
3	15	X	110	Yes	102	82	-67
3	65	Y	18	No	-	0	0
3	120	X & Y	80	No	-	0	0

Determining the winner(s)

Type 2 wins item Y and Type 3 wins item X because the sum of their bids ($130 = 20 + 110$) is greater than the sum of any other bids.

Calculating payments

For Type 1 bidder: She does not win, and therefore pays 0, and gets profit 0.

For Type 2 bidder: **Other winning bid** is 110.

Absent revenue is $128 = 110 + 18$ because bidder Type 3 would have won both items with her individual bids if Type 2's winning bid were absent from the auction.

Type 2's **payment** is $128 - 110 = 18$. The **profit** of Type 2 bidder is $54 - 18 = 36$.

For Type 3 bidder: **Other winning bid** is 20.

Absent revenue is $102 = 82 + 20$ because bidder Type 1 would have won item X and bidder Type 2 would have won item Y if Type 3's winning bid were absent from the auction.

Type 3's **payment** is $102 - 20 = 82$. The **profit** of Type 3 bidder is $15 - 82 = -67$.

Section C

In this section, the Type 3 bidder can only bid on item X and/or item Y individually. She cannot bid on the bundle of items X and Y. Her value for item X or item Y individually is 0.

Type 3 bidder's value for the bundle of items X & Y is a sum of two random numbers between 0 and 100 (with each number equally likely) multiplied by 1.5, but at most 200. Succinctly, it is: $\min([0,100]+[0,100]) \times 1.5, 200$.

Example 7

Bidder type	Value...	...on item(s)	Bid	Winning bid?	Absent Revenue	Payment	Total Profit
1	82	X	17	Yes	40	15	67
2	54	Y	13	No	-	0	0
3	0	X	15	No	-	0	0
3	0	Y	25	Yes	30	13	-13
3	126	X & Y	N/A	-	-	-	-

Determining the winner(s)

Type 1 bidder wins item X and Type 3 bidder wins item Y because the sum of their bids ($42=17 + 25$) is greater than the sum of any other bids.

Calculating payments

For Type 1 bidder: **Other winning bid** is 25.

Absent revenue is $40 = 25 + 15$ because bidder Type 3 would have won both items if Type 1's winning bid were absent from the auction.

Type 1's **payment** is $40 - 25 = 15$. The **profit** of Type 1 bidder is $82 - 15 = 67$.

For Type 2 bidder: She does not win, and therefore pays 0, and gets profit 0.

For Type 3 bidder: **Other winning bid** is 17.

Absent revenue is $30 = 13 + 17$ because bidder Type 1 would have won item X and bidder Type 2 would have won item Y if Type 3's winning bid were absent from the auction.

Type 3's **payment** is $30 - 17 = 13$. The **profit** of Type 3 bidder is $0 - 13 = -13$.

Example 8

Bidder type	Value...	...on item(s)	Bid	Winning bid?	Absent revenue	Payment	Total Profit
1	35	X	30	Yes	49	29	6
2	27	Y	20	Yes	49	19	8
3	0	X	29	No	-	0	0
3	0	Y	19	No	-	0	0
3	78	X & Y	N/A	-	-		

Determining the winner(s)

Type 1 wins item X and Type 2 wins item Y because the sum of their bids ($50 = 30 + 20$) is greater than the sum of any other bids.

Calculating payments

For Type 1 bidder: **Other winning bid** is 20.

Absent revenue is $49 = 20 + 29$ because bidder Type 2 would have won item Y and bidder Type 3 would have won item X if Type 1's winning bid were absent from the auction.

Type 1's **payment** is $49 - 20 = 29$. The **profit** of Type 1 bidder is $35 - 29 = 6$.

For Type 2 bidder: **Other winning bid** is 30.

Absent revenue is $49 = 30 + 19$ because bidder Type 1 would have won item X and bidder Type 3 would have won item Y if Type 2's winning bid were absent from the auction.

Type 2's **payment** is $49 - 30 = 19$. The **profit** of Type 2 bidder is $27 - 19 = 8$.

For Type 3 bidder: She does not win, and therefore pays 0, and gets profit 0.

Example 9

Bidder type	Value...	...on item(s)	Bid	Winning bid?	Absent Revenue	Payment	Total Profit
1	37	X	36	No	-	0	0
2	89	Y	44	No	-	0	0
3	0	X	62	Yes	↓	↓	↓
3	0	Y	50	Yes	↓	↓	↓
3	171	X & Y	N/A	-	80	80	91

Determining the winner(s)

Type 3 bidder wins because the sum of her individual bids on item X and on item Y ($112 = 62 + 50$) is greater than the sum of any other bids.

Calculating payments

For Type 1 bidder: She does not win, and therefore pays 0, and gets profit 0.

For Type 2 bidder: She does not win, and therefore pays 0, and gets profit 0.

For Type 3 bidder: **Absent revenue** is $80 = 36 + 44$ because Type 1 would have won item X and Type 2 would have won item Y if the winning bids of Type 3 were absent from the auction.

Type 3's **payment** is the **absent revenue** of 80. The **profit** of Type 3 is $171 - 80 = 91$.

Section D

In this section, the Type 3 bidder can only bid on the bundle of items X & Y. She cannot bid on item X or item Y individually.

Type 3 bidder's value for the bundle of items X & Y is a sum of two random numbers between 0 and 100 (with each number equally likely) multiplied by 1.5, but at most 200. Succinctly, it is: $\min([0,100]+[0,100]) \times 1.5, 200$.

Example 10

Bidder type	Value...	...on item(s)	Bid	Winning bid?	Absent Revenue	Payment	Total Profit
1	82	X	56	No	-	0	0
2	54	Y	40	No	-	0	0
3	0	X	N/A	-	-	-	-
3	0	Y	N/A	-	-	-	-
3	120	X & Y	105	Yes	96	96	24

Determining the winner(s)

Type 3 bidder wins items X and Y because her bid on the bundle of items X & Y is greater (105) than the sum of the bids of Type 1 and Type 2 bidders (56 + 40 = 96).

Calculating payments

For Type 1 bidder: She does not win, and therefore pays 0, and gets profit 0.

For Type 2 bidder: She does not win, and therefore pays 0, and gets profit 0.

For Type 3 bidder: **Absent revenue** is 96 = 56 + 40 because bidder Type 1 would have won item X and bidder Type 2 would have won item Y with her individual bids if Type 3's winning bid were absent from the auction.

Type 3's **payment** is the **absent revenue** of 96. The **profit** of Type 3 is 120 – 96 = 24.

Example 11

Bidder type	Value...	...on item(s)	Bid	Winning bid?	Absent Revenue	Payment	Total Profit
1	60	X	26	Yes	32	8	52
2	45	Y	24	Yes	32	6	39
3	0	X	N/A	-	-	-	-
3	0	Y	N/A	-	-	-	-
3	48	X & Y	32	No	-	0	0

Determining the winner(s)

Type 1 bidder wins item X and Type 2 bidder wins item Y because the sum of their bids ($50=26 + 24$) is greater than Type 3's bid of 32 on the bundle.

Calculating payments

For Type 1 bidder: **Other winning bid** is 24.

Absent revenue is 32 because bidder Type 3 would have won both items with her bid on the bundle if Type 1's winning bid were absent from the auction.

Type 1's **payment** is $32 - 24 = 8$. The **profit** of Type 1 bidder is $60 - 8 = 52$.

For Type 2 bidder: **Other winning bid** is 26.

Absent revenue is 32 because bidder Type 3 would have won both items with her bid on the bundle if Type 2's winning bid were absent from the auction.

Type 2's **payment** is $32 - 26 = 6$. The **profit** of Type 2 bidder is $45 - 6 = 39$.

For Type 3 bidder: She does not win, and therefore pays 0, and gets profit 0.

Example 12

Bidder type	Value...	...on item(s)	Bid	Winning bid?	Absent Revenue	Payment	Total Profit
1	40	X	39	Tie (No)	-	0	0
2	70	Y	69	Tie (No)	-	0	0
3	0	X	N/A	-	-	-	-
3	0	Y	N/A	-	-	-	-
3	141	X & Y	108	Tie (Yes)	108	108	33

Determining the winner(s)

There is a **tie** since the sum of Type 1 bidder's bid and Type 2 bidder's bid is exactly equal to Type 3 bidder's bid. The computer simulates a coin toss. Type 3 wins.

Calculating payments

For Type 1 bidder: She does not win, and therefore pays 0, and gets profit 0.

For Type 2 bidder: She does not win, and therefore pays 0, and gets profit 0.

For Type 3 bidder: **Absent revenue** is $108 = 39 + 69$ because Type 1 would have won item X and Type 2 would have won item Y if the winning bid were absent from the auction.

Type 3's **payment** is the **absent revenue** of 108. The **profit** of Type 3 is $141 - 108 = 33$.

Solution to the Understanding Test

Table 1

Bidder type	Value...	...on item(s)	Bid	Winning bid?	Absent Revenue	Payment	Total Profit
1	36	X	26	Yes	80	14	22
2	70	Y	66	Yes	54	28	42
3	35	X	14				
3	65	Y	28				
3	150	X & Y	N/A				

Table 2

Bidder type	Value...	...on item(s)	Bid	Winning bid?	Absent Revenue	Payment	Total Profit
1	36	X	31				
2	70	Y	33				
3	35	X	34	Yes	↓	↓	↓
3	65	Y	90	Yes	↓	↓	↓
3	150	X & Y	N/A		64	64	86

Table 3

Bidder type	Value...	...on item(s)	Bid	Winning bid?	Absent Revenue	Payment	Total Profit
1	32	X	25	Yes	65	20	12
2	70	Y	16				
3	35	X	20				
3	65	Y	45	Yes	41	16	49
3	150	X & Y	N/A				

Table 4

Bidder type	Value...	...on item(s)	Bid	Winning bid?	Absent Revenue	Payment	Total Profit
1	87	X	85	Yes	105	25	62
2	90	Y	47				
3	10	X	25				
3	50	Y	80	Yes	132	47	3
3	90	X & Y	N/A				

References

- Ansari, A. R. and R. A. Bradley (1960). Rank-sum tests for dispersions. *Annals of Mathematical Statistics* 31(4), 1174–1189.
- Ausubel, L. M. and O. V. Baranov (2010, August). Core-selecting auctions with incomplete information. Mimeo.
- Beck, M. and M. Ott (2011). Revenue monotonicity in core-selecting package auctions. Mimeo.
- Brunner, C., J. K. Goeree, C. A. Holt, and J. O. Ledyard (2010). An experimental test of flexible combinatorial spectrum auction formats. *American Economic Journal: Microeconomics* 2(1), 39–57.
- Bykowsky, M., R. Cull, and J. Ledyard (2000). Mutually destructive bidding: The FCC auction design problem. *Journal of Regulatory Economics* 17(3), 205–228.
- Chen, Y. and K. Takeuchi (2010). Multi-object auctions with package bidding: An experimental comparison of Vickrey and iBEA. *Games and Economic Behavior* 68(2), 557–579.
- Chernomaz, K. and D. Levin (2012). Efficiency and synergy in a multi-unit auction with and without package bidding: An experimental study. *Games and Economic Behavior* 76(2), 611–635.
- Costa-Gomes, M. A. and V. P. Crawford (2006). Cognition and behavior in two-person guessing games: An experimental study. *American Economic Review* 96(5), 1737–1768.
- Cox, J. C., V. L. Smith, and J. M. Walker (1988). Theory and individual behavior of first-price auctions. *Journal of Risk and Uncertainty* 1, 61–99.
- Cramton, P. C., Y. Shoham, and R. Steinberg (2006). *Combinatorial Auctions*. MIT Press.
- Crawford, V. P. and N. Iriberri (2007). Level-k auctions: Can a nonequilibrium model of strategic thinking explain the winner’s curse and overbidding in private-value auctions? *Econometrica* 75(6), 1721–1770.
- Crawford, V. P., T. Kugler, Z. Neeman, and A. Pauzner (2009). Behaviourally optimal auction design: Examples and observations. *Journal of the European Economic Association* 7(2-3), 377–387.
- Day, R. W. and P. C. Cramton (2008). Quadratic core-selecting payment rules for combinatorial auctions. Working paper, University of Maryland.

- Day, R. W. and P. C. Cramton (2012). Quadratic core-selecting payment rules for combinatorial auctions. *Operations Research* 60(3), 588–603.
- Day, R. W. and P. Milgrom (2008). Core-selecting package auctions. *International Journal of Game Theory* 36(3-4), 393–407.
- Day, R. W. and S. Raghavan (2007). Fair payments for efficient allocations in public sector combinatorial auctions. *Management Science* 53(9), 1389–1406.
- Dyer, D., J. H. Kagel, and D. Levin (1989). A comparison of naive and experienced bidders in common value offer auctions: A laboratory analysis. *Economic Journal* 99(394), 108–115.
- Epstein, R., L. Henríquez, J. Catalán, G. Y. Weintraub, C. Martínez, and F. Espejo (2004). A combinatorial auction improves school meals in chile: a case of OR in developing countries. *International Transactions in Operational Research* 11(6), 593–612.
- Erdil, A. and P. Klemperer (2010). A new payment rule for core-selecting package auctions. *Journal of the European Economic Association* 8(2-3), 537–547.
- Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics* 10(2), 171–178.
- Gale, I. (1990). A multiple-object auction with superadditive values. *Economics Letters* 34(4), 323–328.
- Goeree, J. K. and C. A. Holt (2010). Hierarchical package bidding: A paper & pencil combinatorial auction. *Games and Economic Behavior* 70(1), 146–169.
- Goeree, J. K. and Y. Lien (2009, September). An equilibrium analysis of the simultaneous ascending auction. Working Paper 428, Institute for Empirical Research in Economics University of Zurich.
- Goeree, J. K. and Y. Lien (2012). On the impossibility of core-selecting auctions. Mimeo.
- Hafalir, I. E. and H. Yektaş (2012, February). Core deviation minimizing auctions. Mimeo.
- Jehiel, P. (2011). Manipulative auction design. *Theoretical Economics* 6(2), 185–217.
- Kagel, J. H. and D. Levin (2005). Multi-unit demand auctions with synergies: Behavior in sealed-bid versus ascending-bid uniform-price auctions. *Games and Economic Behavior* 53(2), 170–207.

- Kagel, J. H. and D. Levin (2011). *Handbook of Experimental Economics*, Volume 2, Chapter Auctions: A Survey of Experimental Research, 1995 – 2010. Princeton University Press.
- Kagel, J. H., Y. Lien, and J. K. Goeree (2012, March). Ascending prices and package bidding: Further experimental analysis. Technical report.
- Kagel, J. H., Y. Lien, and P. Milgrom (2010). Ascending prices and package bidding: A theoretical and experimental analysis. *American Economic Journal: Microeconomics* 2(3), 160–85.
- Kelso, A. S. and V. P. Crawford (1982). Job matching, coalition formation, and gross substitutes. *Econometrica* 50(6), 1483–1504.
- Klemperer, P. (2004). *Auctions: Theory and Practice*. Princeton University Press.
- Knetsch, J. L., F.-F. Tang, and R. H. Thaler (2001). The endowment effect and repeated market trials: Is the Vickrey auction demand revealing? *Experimental Economics* 4(3), 257–269.
- Krishna, V. (2009). *Auction Theory*. Academic Press.
- Krishna, V. and R. Rosenthal (1996). Simultaneous auctions with synergies. *Games and Economic Behavior* 17(1), 1–31.
- Kwasnica, A., J. Ledyard, D. Porter, and C. DeMartini (2005). A new and improved design for multiobject iterative auctions. *Management Science* 51(3), 419–34.
- Lamy, L. (2010). Core-selecting package auctions: a comment on revenue-monotonicity. *International Journal of Game Theory* 39(3), 503–510.
- Ledyard, J. O., D. Porter, and A. Rangel (1997). Experiments testing multiobject allocation mechanisms. *Journal of Economics & Management Strategy* 6(3), 639–675.
- Levin, J. (1997). An optimal auction for complements. *Games and Economic Behavior* 18(2), 176–192.
- Lunander, A. and S. Lundberg (2012). Bids and costs in combinatorial and noncombinatorial procurement auctions – evidence from procurement of public cleaning contracts. *Contemporary Economic Policy*.
- Marszalec, D. (2011). *Essays on Auctions*. Ph. D. thesis, Nuffield College, University of Oxford.
- Maskin, E. and J. Riley (2000). Asymmetric auctions. *Review of Economic Studies* 67(3), 413–438.

- McMillan, J. (1994). Selling spectrum rights. *Journal of Economic Perspectives* 8(3), 145–162.
- Milgrom, P. (2000). Putting auction theory to work: The simultaneous ascending auction. *Journal of Political Economy* 108(2), 245–272.
- Milgrom, P. (2004). *Putting Auction Theory to Work*. Cambridge University Press.
- Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research* 6(1), 58–73.
- Neugebauer, T. and R. Selten (2006). Individual behavior of first-price auctions: The importance of information feedback in computerized experimental markets. *Games and Economic Behavior* 54(1), 183 – 204.
- Ockenfels, A. and R. Selten (2005). Impulse balance equilibrium and feedback in first-price auctions. *Games and Economic Behavior* 51(1), 155 – 170.
- Plott, C. R. (1997). Laboratory experimental testbeds: Application to the PCS auction. *Journal of Economics & Management Strategy* 6(3), 605–638.
- Roth, A. E. (2008). What have we learned from market design? *Economic Journal* 118(527), 285–310.
- Sano, R. (2011). Incentives in core-selecting auctions with single-minded bidders. *Games and Economic Behavior* 72(2), 602–606.
- Sheffi, Y. (2004). Combinatorial auctions in the procurement of transportation services. *Interfaces* 34(4), 245–252.
- Shi, X. (2013, July). A nondegenerate Vuong test with application to modeling voter turnout. Technical report, University of Wisconsin - Madison.
- Sun, N. and Z. Yang (2006). Equilibria and indivisibilities: Gross substitutes and complements. *Econometrica* 74(5), 1385–1402.
- Sun, N. and Z. Yang (2008). A double-track auction for substitutes and complements. Discussion Paper 656, Kyoto Institute for Economic Research.
- Sun, N. and Z. Yang (2009). A double-track adjustment process for discrete markets with substitutes and complements. *Econometrica* 77(3), 933–952.
- Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* 16(1), 8–37.
- Vuong, Q. H. (1989). Likelihood ratio tests for model selection and non-nested hypotheses. *Econometrica* 57(2), 307–333.

Chapter 5

Friending: a model of online social networks

Abstract

We develop a parsimonious and tractable dynamic social network formation model in which agents interact in overlapping social groups. The model allows us to analyse network properties and homophily patterns simultaneously. We derive closed-form analytical expressions for the distributions of degree and, importantly, of homophily indices using mean-field approximations. We test the comparative static predictions of our model using a large dataset from Facebook covering student friendship networks in ten American colleges in 2005, and calibrate the analytical solutions to these networks. We find good empirical support for our predictions. Furthermore, at the best-fitting parameter values, the homophily patterns, degree distribution, and individual clustering coefficients resulting from the simulations of our model fit well with the data. Our best-fitting parameter values indicate how American college students allocate their time across various activities when socialising.

Keywords: social networks, homophily, network formation

JEL Classification: *C63, D85, Z13*

And what so poor a man as Hamlet is
May do, to express his love and friending to you,
God willing, shall not lack.

(William Shakespeare, *Hamlet*, Act 1, Scene 5)

5.1 Introduction

Friendships are an essential part of economic life and social networks affect many areas of public policy.¹ Friendships create externalities, which impact educational performance (Sacerdote, 2001), health (Kremer and Levy, 2008), group lending (Banerjee et al., 2012), and productivity at work (Falk and Ichino, 2006). Recently, online social networks have become a global record of naturally occurring social ties. The world’s largest online social network – Facebook – is increasingly becoming the main platform for interacting with friends and documenting friendships.² Launched in 2004 and at first exclusive to American colleges, it now has over a billion active users worldwide.³ An average user spends 405 minutes on Facebook per month.⁴ Facebook allows users to share pictures, videos, links, as well as organise events, play games, and develop professional contacts through numerous third-party applications. On Facebook, users have access to a huge amount of information about other users, which influences the network formation process (Lewis et al., 2008, 2012). In this chapter, we propose a social network formation model, which uses this information to explain who befriends whom on Facebook.

5.1.1 Homophily

A particular focus of this chapter is *homophily* – the tendency of individuals to associate with those who are similar to themselves – which has been well documented in

¹The best recent summaries of applications of networks in the social sciences are by Jackson (2008), Goyal (2009), Easley and Kleinberg (2010), and Newman (2010).

²Since 2011 Facebook has become the dominant online social network in almost every country in the world except China, Russia, Belarus, Ukraine, Iran, Armenia, Kazakhstan, Latvia, and Vietnam.

³Active users are those who logged on to their Facebook profile at least once in the previous month. See SEC Form 10-Q 2012Q2 filing: <http://www.sec.gov/Archives/edgar/data/1326801/000119312512325997/d371464d10q.htm>.

⁴This is far more than on any other social networking website: on average users spent 17 minutes on LinkedIn, 8 minutes on MySpace and 21 minutes on Twitter and 3 minutes on Google+ per month. These data come from a Bloomberg report based on a comScore study: <http://www.bloomberg.com/news/2012-02-28/google-users-spent-less-time-on-site-in-january-comscore-finds.html>. Pempek et al. (2009) found similar Facebook use intensity for college students.

sociology.⁵ Homophily patterns, for example, play an important role in school segregation (Currarini et al., 2009) and information transmission (Golub and Jackson, 2012). There are also many studies regarding the causes of homophily. Some empirical studies in economics (Mayer and Puller, 2008) and sociology (Moody, 2001, Mouw and Entwisle, 2006) find that most of the observed homophily can be explained by a bias in people’s preferences. More recently, Currarini et al. (2009, 2010) proposed a rigorous model explaining several striking patterns of homophily in ethnicity in high-school peer groups. Yet Currarini et al. (2009) make it clear that the observed racial homophily patterns do not necessarily arise from an exogenous bias in preferences towards people of the same type. Rather, similar people may be simply more likely to meet. Wimmer and Lewis (2010) provide some support for that idea by studying racial homophily in a small Facebook dataset. They find that sharing the same physical environment⁶ and reciprocal friendships, are far more important in explaining homophily than race preference.

5.1.2 Socialising on Facebook

In social networks, the characteristics of the agents constitute the *identity* of the person they represent. As Sen (2006) emphasises, a person’s identity is necessarily multidimensional: one can simultaneously identify oneself as a woman, a student, a Catholic, a vegetarian, and a rower. An identity is then a collection of characteristics drawn from *social categories*.⁷ In the preceding example, the social categories are: gender, employment status, religion, dietary practice, and sport activity. A *social group* is a collection of persons sharing a characteristic from a particular social category.

Let us immediately make these ideas more concrete and think about two students

⁵The two classic studies of homophily in humans by Kandel (1978) and Shrum et al. (1988) found racial and gender homophily in adolescent social groups. McPherson et al. (2001) provide an excellent survey of the literature and cite numerous examples of homophily.

⁶The authors call this “propinquity”.

⁷Akerlof and Kranton (2010) summarise the importance of identity in economics.

– Mark and Eduardo – who are “friends” on the Facebook network of a prestigious American university. Mark and Eduardo live in the same dorm, but Mark is a computer science major, whereas Eduardo studies finance. There are many processes that explain how Mark and Eduardo became friends on Facebook. In our model, we propose that Mark and Eduardo allocate time across their various social categories, such as attending lectures and classes and spending time in their dorm. Naturally, a lot of the time allocation is determined institutionally by timetables or geographical locations. The overlap between their social groups (and their relative sizes) determines how frequently they interact with each other socially and their chance of meeting in person. If Mark and Eduardo were also members of the same fraternity, their chance of meeting would be even higher. Their eventual friendship is then documented online via Facebook.

5.1.3 Our contribution

This chapter makes several contributions. We develop an intuitive and parsimonious dynamic social network formation model. The process governing friendship formation resembles our description of how Mark and Eduardo become friends on Facebook. That is, agents allocate time across various social categories thus determining how frequently they interact with others in each social group. When interacting with others in a social group, an agent forms a friendship with another agent chosen at random from among those in the group, who are not yet his friends and who are still actively using Facebook.

We are interested in the structural properties of the resulting network. We are able to obtain closed-form analytical expressions for the degree distribution and for various measures of homophily. Importantly, we are able to derive the full distribution of individual homophily indices. We calibrate the parameters in these expressions against the Facebook data.

The entire process is governed by the allocation of time and the relative sizes of the groups to which the agents belong. Since agents with certain sets of characteristics may

interact more often, homophily may emerge with respect to particular social categories. As such, the biases in the frequency of interaction between agents in our model can either be seen as a pure bias in meeting opportunities, or as the manifestation of agents' preferences over how they allocate their time. Our model, therefore, does not distinguish these two possible effects. Furthermore, since choices are made stochastically, we bypass strategic considerations for friendship formation. However, this simplification allows us to develop a dynamic network formation model in which agents' characteristics determine the formation process.

In this chapter, we focus on homophily for *immutable* social categories, such as gender or – in the context of a university – year of graduation, because there is no feedback mechanism that would allow agents to change their characteristics within these social categories on the basis of their friendships.

The empirical part of this chapter provides striking support for our model. We find the best-fitting parameter values, which determine the allocation of time across social categories, and best fit the degree and homophily distributions in gender and year of graduation for ten separate Facebook networks. Students' friendships reveal that they spend more time socialising in class than in their dorms. The model fits the data extremely well despite its parsimony (there are three degrees of freedom). Remarkably, the simulations run at the best-fitting parameter values show that the individual clustering coefficient distributions also match the clustering patterns in all the networks.

Following a brief literature review in the next section, the outline of this chapter is as follows. In section 5.3, we formally present and discuss the social network formation model. In section 5.4, we derive the degree distribution and homophily indices (as well as other properties of the network) using a mean-field approximation method and test this approximation in section 5.5 against simulation results. Sections 5.6 and 5.7

present the Facebook dataset and explore some baseline empirical patterns. In section 5.8, we calibrate the model to the data, and present our empirical results. Section 5.9 discusses four possible extensions to the model and section 5.10 concludes. The Appendix contains the algorithms and methods used for calibration, a full description of the data, the results table, and further empirical tests of the model.

5.2 Related literature

In many dynamic social network formation models, agents (represented by nodes in a graph) are anonymous. The formation of new friendships (edges or links) then depends entirely on the existing links in the network. In a seminal paper, Barabási and Albert (1999) proposed a model in which every node receives a link with a probability proportional to its existing number of links. In this *preferential attachment* framework, Mark would be more likely to send a “friend request” to Eduardo if the latter is already popular. We discuss this approach further in section 5.9.3. Jackson and Rogers (2007) additionally suggested that “friends of friends” are more likely to link. Hence, if Mark knows that he and Eduardo have a Facebook friend in common, then he and Eduardo are likely to establish a direct Facebook link with each other. These types of models provide analytical expressions and comparative statics for the macroscopic properties of the network: degree distribution, clustering, diameter, average distance, and assortativity.

However, these models are unable to explain homophily patterns, since node characteristics are not taken into account. Node characteristics can play a big role in explaining the topology of a network (Bianconi et al., 2009). One branch of the economics literature explores the equilibria and stability of static networks, where node characteristics determine the linking process (de Marti and Zenou, 2011, Iijima and Kamada, 2013). We contribute to another branch, which considers dynamic processes. Currarini et al. (2009) originally proposed a dynamic matching model with a biased

meeting process in which agents prefer to link to those who are similar to themselves. Agents were endowed with a characteristic from one social category, and the biased meeting process is determined by an exogenous parameter. Given the nature of the model, it cannot account for the properties of the resulting network of friendships.⁸ Bramoullé et al. (2012) extended the model of Jackson and Rogers (2007) to consider homophily in a random growing network with multidimensional node characteristics and tested the comparative static predictions of their model against a dataset of empirical citation networks. Our approach is similar in spirit to their paper, but complements it in several important ways. As in our model, networks form through the formation of new links over time. However, most of the results given in Bramoullé et al. (2012) are in the limit as time approaches infinity. Some results are also given for *any* time period for the case in which there are only two relevant social groups. In contrast, our model offers new theoretical results for *any* time period. In particular, we replicate the result Bramoullé et al. (2012) in the case of two social groups (showing that homophily becomes a decreasing function of time and degree), but we also show how this result breaks down in the case of multiple social groups. In addition, although Bramoullé et al. (2012) are able to derive properties of the resulting networks and obtain comparative statics on the relationship between homophily and degree, they do not obtain a closed-form solution for the degree distribution. Since we are able to obtain closed-form solutions for our expressions, we are able not only to test the comparative static predictions of our model, but also to calibrate the model to our dataset, thereby isolating the best-fitting parameter values of our model. We are not aware of any studies that carry out such a calibration. Finally, in fitting the model to the data, we consider entire distributions of degree, homophily indices, and clustering coefficients rather than

⁸They use data from *The National Longitudinal Study of Adolescent Health* (Add Health), which represents a relatively restricted network structure. Students were asked to name their ten “best friends” and around three quarters of students choose to nominate fewer than ten “best friends.” Additionally, at most 5 of them could be of the same sex. This means that a deep analysis of the network properties is not usually possible.

simply fitting averages; this is also something we have not encountered in the literature.

Our model is also conceptually related to *affiliation networks* introduced in sociology by Breiger (1974) and Feld (1981). We discuss this relationship further in sections 5.3.5 and 5.9.4. An affiliation network is described by a set of agents and a set of *memberships*, such as clubs, online fora, research topics, or social groups (Newman et al., 2002). These models have found wide-spread application in online social networks (Botha and Kroon, 2010, Kumar et al., 2010, Xiang et al., 2010). In more recent evolving models of affiliation networks, new memberships may emerge over time, and the likelihood of meeting new agents can depend on their memberships (Lattanzi and Sivakumar, 2009, Zheleva et al., 2009). However, these models typically contain a large number of parameters and most, such as those by Leskovec et al. (2005) and Leskovec et al. (2008), rely entirely on simulations.

5.3 Model

5.3.1 Characteristics of agents

Let $\mathcal{K} = [K^0, \dots, K^R]$ be a finite ordered list of *social categories*. An element K^r is the r^{th} category and $k \in K^r$ is a *characteristic* within that category. Let $\mathcal{R} = \{0, \dots, R\}$ and $\mathcal{R}_+ = \mathcal{R} \setminus \{0\}$. The *identity* of every agent $i \in N$ is represented by a vector $\mathbf{k}_i = (k_i^0, \dots, k_i^R)$ of characteristics, where for each $r \in \mathcal{R}$, $k_i^r \in K^r$.⁹ For any pair $i, j \in N$, let $k_i^0 = k_j^0$.¹⁰ For each $r \in \mathcal{R}$, define a *social group* $\Gamma_i^r = \{j \in N \mid k_i^r = k_j^r\} \setminus \{i\}$, which is the set of all agents (other than i) that share the characteristic k_i^r within the social category r with i . Note that $\Gamma_i^0 = N \setminus \{i\}$. Finally, for each non-empty $S \subseteq \mathcal{R}$,

⁹In effect, we assume that characteristics are mutually exclusive within a particular social category, such as race or gender. The model can be extended to allow agents to have subsets of characteristics (i.e. $k_i^r \subseteq K^r$), such as classes within a particular major.

¹⁰This does not restrict the characteristics space in any way. The zeroth category, which greatly simplifies notation, is one in which all agents share the same characteristic.

define

$$\pi_i(S) = \bigcap_{r \in S} \Gamma_i^r \setminus \bigcup_{r \in \mathcal{R} \setminus (S \cup \{0\})} \Gamma_i^r \quad (5.1)$$

which induces a partition $\Pi_i = \{\pi_i(S) | S \subseteq \mathcal{R}, S \neq \emptyset\}$ on $N \setminus \{i\}$.¹¹ Therefore, $\pi_i(S)$ is the set of agents (other than i) that share *only* the characteristics within the set of categories indexed by S with i .

Example. Consider an online social network at a university in which we can observe the class, dorm, gender, and year of graduation for each student. Then using our notation

$$\mathcal{K} = [K^0, K^1, K^2, K^3, K^4] = [\textit{student}, \textit{class}, \textit{dorm}, \textit{gender}, \textit{year of graduation}]$$

All agents are students ($k_i^0 = k_j^0$ for all $i, j \in N$). Thus $K^1 \in \mathcal{K}$, which represents class, can include $k \in \{\textit{maths}, \textit{computer science}, \textit{psychology}\}$. Suppose, that the identity of a student i is represented by a vector

$$\mathbf{k}_i = (\textit{student}, \textit{computer science}, \textit{Kirkland House}, \textit{male}, 2006)$$

Let us consider $S = \{1, 3\}$. Now, i 's social group Γ_i^1 is the set of all computer science students and Γ_i^3 is the set of all male students (other than i). Then, $\pi_i(S)$ is the set of male students other than i , who take the computer science class, but do not share any other characteristics with i . $\pi_i(\{0\})$ would be the set of all female non-computer-scientists, who do not live in Kirkland House and are not graduating in 2006. Π_i represents the partition into disjoint sets of students, who share exactly 1, 2, 3, 4 or 5 social categories with i . ■

¹¹Note that $\pi_i(S) = \pi_i(S \cup \{0\})$ for all non-empty $S \subseteq \mathcal{R}$. Furthermore, since $\Gamma_i^r = \bigcup_{\pi \in \{\pi_i(S) | r \in S\}} \pi$, a social group is a union of disjoint partition elements.

5.3.2 Network formation process

We model our network as a simple, undirected graph, with a finite set of nodes (which represent agents) and a finite set of edges (which represent friendships). The *degree* of an agent i is the number of i 's friends. At time period $t = 0$ all agents are *active* and have no friends. Let $\mathbf{q} = (q^0, \dots, q^R)$ and $\sum_{r \in \mathcal{R}} q^r = 1$. In each period $t \in \{1, 2, 3 \dots\}$, an active agent interacts with agents in the social group Γ_i^r with probability $q^r \geq 0$. One can think of $\Gamma_i^0 = N \setminus \{i\}$ as the social group that i interacts with during i 's "free time". We can thus interpret q^r as the proportion of time in any period t that agent i spends with agents in the social group Γ_i^r . During the interaction in a social group, an agent i is linked to another active agent in that group chosen uniformly at random with whom i is not yet a friend. If the agent is already linked to every other active agent in that social group, the agent makes no friends in that period. Friendships are always reciprocal so all links are undirected. Finally, in every period, agent i remains *active* with a given probability $p_i \in (0, 1)$ until the following period and becomes *idle* with probability $1 - p_i$. If the agent i becomes idle, i retains all her friendships, but can no longer form any links with other agents in all subsequent periods.¹²

5.3.3 Interpretation of the model

We can interpret our model in the context of an online social network, such as Facebook. We imagine that the online social network has users, who interact with each other either online or offline. Users can meet each other physically in real-world social groups: for example, university students could meet in class, in their dorm, or at parties during their free time. This was particularly relevant in the earlier stages of Facebook when it was only open to selected American colleges. Additionally, most social networks

¹²Naturally, i cannot make a link to oneself.

allow users to browse profiles of other users according their memberships in particular social groups (Facebook, for example, facilitates direct browsing of users' profiles by characteristic).

In our model, \mathbf{q} could represent the fraction of time that students physically spend within various social categories or their propensity to browse for other students of these social categories online. We assume that, although Mark and Eduardo may spend the same time in class (since \mathbf{q} is the same for both of them), they will be meeting different social groups of people if Mark is attending a computer science lecture and Eduardo is taking a finance course. When Mark is interacting with other computer scientists, he befriends them (at random) and then documents these friendships via the online social network (for example, by sending a “friend request” on Facebook).¹³ Even after every computer scientist in Mark's lecture becomes his friend, Mark still attends the lecture. Henceforth, whenever Mark spends time in the class, he does not make any more friends with the lecture attendees. However, Mark could still be making friends with students in his other social groups: for example, with finance majors in his fraternity.

Social categories are, technically, just the names of variables, which we can observe about the users of an online social network. Since these social categories can be virtually anything, it does not always make sense to impose a positive probability q^r of spending time in every social category r . For example, gender and graduation year can be social categories, but it is not very meaningful to say that Mark *specifically* allocates time to spend it only with men or only with students of his graduation year. Rather, Mark may be more or less likely to meet these students because of the classes he takes or the dorm he lives in. This point will be relevant when we fit the model to the data and we return to it in section 5.8.

There are several ways of interpreting $1 - p_i$, the probability of becoming idle. There

¹³Our model assumes that all “friend requests” are accepted.

must be reasons, *other than having linked with every user in the network*, for why people stop adding new friends online: losing interest, finding an alternative online social network, reaching a cognitive capacity for social interaction, and so on. Including all these explanations would require a much richer model so we simply capture them as a random process with the idleness probability $1 - p_i$. One is to imagine that it represents the probability that in any period Eduardo stops sending or accepting ‘friend requests’ even though he may still be actively using the online platform to stay in touch with his current friends.¹⁴

Example (cont.) Figure 5.1 succinctly summarises the link formation process in our example and its interpretation for agent i . This process happens simultaneously for all agents in every period. Furthermore it is assumed that $q^3 = q^4 = 0$. ■

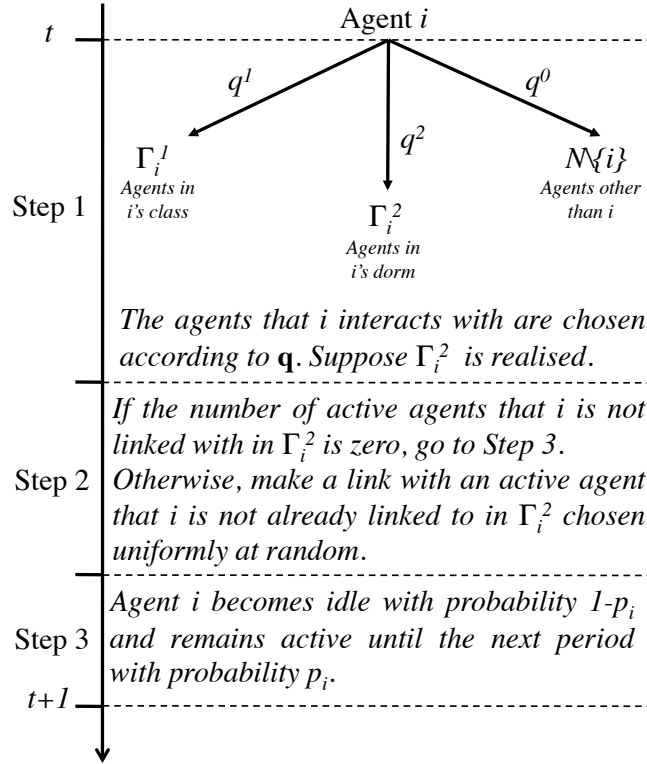
5.3.4 Discussion of the model

Many dynamic social network formation models are growing random network models in which new nodes arrive in every period and link to existing ones (e.g. Price (1976), Barabási and Albert (1999), Jackson and Rogers (2007)). In contrast, we chose to present a model with a fixed number of nodes, which become idle while retaining their links. This choice is not intellectual curiosity alone. One can, of course, think of Facebook as a growing network (new users join every day); however, our model allows us to focus on the formation of links among the existing users. Naturally, it is possible to extend our model to accommodate for the arrival of new users, and we discuss this in section 5.9.1.

Our model has a unique theoretical feature. For every agent in every social group, we derive an *expected stopping time* at which the agent makes friends with every active agent in that group. This highlights the idea that all interactions within social groups

¹⁴We think of fixed $1 - p_i$ as a crude approximation to a function which represents agents’ likelihood of becoming idle. Estimating such a function could be a fruitful area for further research. See a further discussion in section 5.9.2.

Figure 5.1: Network formation process in the Example



are inherently local. Yet we are able to characterise the macroscopic properties of the network in terms of these expected stopping times and $(p_i)_{i \in N}$ alone.

We have set up our model in a manner that does not require agents to make any optimal decisions. Agents do not maximise a utility function, but rather all their choices are fully stochastic. We could have similarly assumed that observed friendship choices resulted from optimal decisions of utility-maximising agents (Currarini et al. (2009) take this approach). Indeed, endow every agent $i \in N$ with a utility function $U_i(d_i) = v_i(d_i) - c_i d_i$ where d_i is the number of i 's friends and c_i is the marginal cost of creating a new friendship. Suppose that, in every period, every agent “spends time” in some social group, and can choose one “active” agent within that social group with whom she makes a new friendship. Since the characteristics of that agent are irrelevant to i 's utility, the specific agent that is chosen does not matter and so he can

be chosen uniformly at random. If we also assume that the benefit function $v_i(\cdot)$ is strictly increasing, twice-differentiable and concave in its argument, then there will be a finite number of friends d_i^* satisfying $v_i'(d_i^*) = c_i$.¹⁵ Agent i will therefore keep adding friends in every period up to the point at which she has d_i^* friends. What remains is to find a family of utility functions $\{U_i(\cdot)\}_{i \in N}$ such that the distribution of d_i^* matches $G(d)$ in equation (5.13) below, and we obtain a model equivalent to the one we outlined in section 3.2.

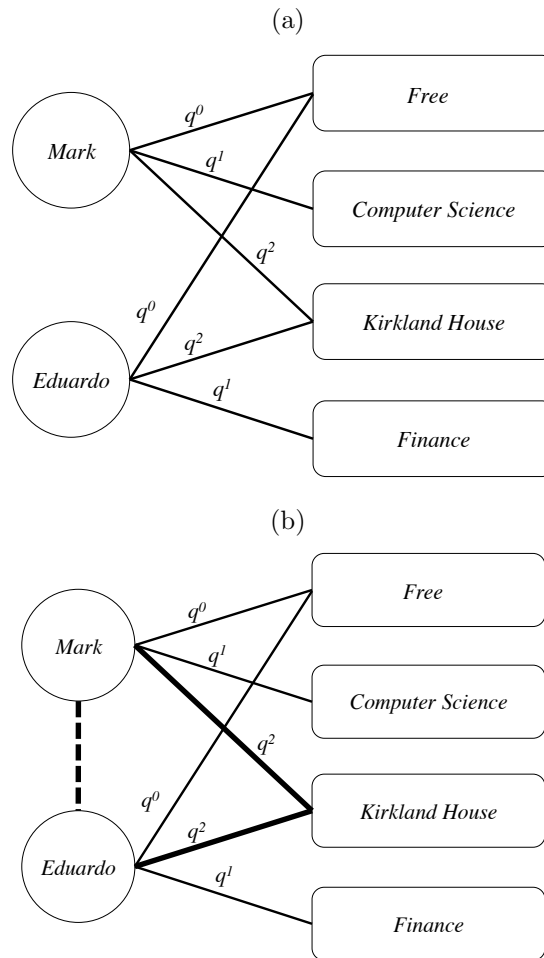
5.3.5 Relationship to affiliation networks

Our network formation process can be reinterpreted as a dynamic *affiliation network*. An affiliation network is initialised by a bipartite graph consisting of two sets of nodes: agents and memberships. In our framework, memberships correspond to characteristics, for example, computer science class or Kirkland House dorm. At the beginning, the only links in this graph are between agents and memberships. New links can be formed by “closing” transitive triples: if two agents are linked to the same membership, there is a positive probability that a link will form between them. Easley and Kleinberg (2010, p. 97) call this *focal closure*.

More specifically within our framework, the set of all memberships is $\{k \in K^r | r \in \mathcal{R}\}$ and a link between an agent i and a membership $k \in K^r$ is given the strength q^r for all $i \in N$. Figure 5.2(a) represents our Example as an initialised bipartite graph of an affiliation network. The formation of new links via focal closure happens in the following way: in every period, every agent i is assigned a membership $k \in K^r$ at random according to \mathbf{q} , and the agent forms a link with another agent j chosen uniformly at random from among the remaining active agents that have a link with $k \in K^r$. For example, Mark and Eduardo could become friends in some period because Mark was assigned to Kirkland House with which Eduardo also has a membership. This

¹⁵Ignoring the slight complication that d_i^* must be an integer.

Figure 5.2: Model as an affiliation network



is shown in Figure 5.2(b).

5.4 Theoretical results

We are interested in analysing properties of the network generated by the model. In order to derive closed-form expressions for the degree distribution and homophily indices, we use the mean-field approximation method used in statistical mechanics (see Barabási et al. (1999), Jackson and Rogers (2007)). We assume that the realisation of a random variable in any period is its expected value. Hence, the dynamic system generated by our model does not evolve stochastically, but rather deterministically

at the rate proportional to the expected change. Mean-field approximation has been adopted by economics literature, and our analysis is similar to the one carried out by Jackson and Rogers (2007). In general, the mean-field approximation method is not without its drawbacks. The accuracy of its predictions must be tested against simulations (Jackson, 2008, p. 137). In section 5.5, we show that the approximation works well for our model.

5.4.1 Degree distribution

In order to derive the degree distribution, we first analyse the meeting process of agents across social groups. The probability with which agent i interacts with an agent from $\pi_i(S)$ is given by

$$q^{\pi_i(S)} = |\pi_i(S)| \left[\sum_{r \in S \cup \{0\}} \frac{q^r}{|\Gamma_i^r|} \right] \quad (5.2)$$

and by definition $\sum_{\pi \in \Pi_i} q^\pi = 1$.

Example (cont.) To understand equation (5.2), let us derive $q^{\pi_i(\{1\})}$ in our example. This is the proportion of time that i spends with students that are in his class, but not in his dorm. There are $|\Gamma_i^1|$ students in i 's class in total, and there are $|\pi_i(\{1\})|$ who are both in his class but not in his dorm. He can encounter students in his class but not in his dorm either during the time he spends in class or during his free time. When in class, which happens with probability q^1 , he encounters students who are in his class but not in his dorm with probability $\frac{|\pi_i(\{1\})|}{|\Gamma_i^1|}$. Similarly, during his free time, which happens with probability q^0 , he encounters students who are in his class but not in his dorm with probability $\frac{|\pi_i(\{1\})|}{|N|-1}$. Hence, the proportion of time that i spends with students that are in his class, and not in his dorm is given by $q^{\pi_i(\{1\})} = |\pi_i(\{1\})| \left[\frac{q^1}{|\Gamma_i^1|} + \frac{q^0}{|N|-1} \right]$. ■

Let $d_i(t)$ be the degree of agent i in period t . Analogously, $d_i^\pi(t)$ is the number of friends i has in period t with agents in $\pi \in \Pi_i$. If T^π is the expected time it takes i to make a link with every other active agent in π (expected stopping time), then the

mean-field approximation of the degree change of i with agents in π between periods t and $t + 1$ is

$$\Delta d_i^\pi(t) = q^\pi \left(1 + \frac{R^\pi(t)}{R^\pi(t)} \right) \mathbf{1}(t \leq T^\pi) = 2q^\pi \times \mathbf{1}(t \leq T^\pi) \quad (5.3)$$

where $R^\pi(t)$ is the total number of remaining active agents in π (other than i) with whom i is not yet linked at time t and $\mathbf{1}$ is an indicator function. In other words, conditional on being in π , i makes a link to an agent in π . Agent i also receives one link on average from an agent in π at t : there are $R^\pi(t)$ other active agents (with whom i is not linked) in π , and each is linked with i with probability $\frac{1}{R^\pi(t)}$.¹⁶ Hence, “on average” i makes $2q^\pi$ friends in π in every period until T^π .

The partition Π_i induced on $N \setminus \{i\}$ allows us to consider the links made between agents of any element π separately. Hence, we determine the link formation process within each element $\pi \in \Pi_i$ and then weigh it by q^π – the proportion of time spent in π . Despite this analytical trick, the actual network formation process certainly allows agents to receive links from outside the social group they are currently interacting in. This fact also justifies our ignoring the possibility that any two agents make the same link simultaneously, which is negligible for large N .

Recall that, in period $t = 0$, every agent i has no friends. Solving equation (5.3) with our initial condition $d_i^\pi(0) = 0$ gives

$$d_i^\pi(t) = 2q^\pi [t\mathbf{1}(t \leq T^\pi) + T^\pi\mathbf{1}(t > T^\pi)] \quad (5.4)$$

In order to obtain the expected stopping time T^π for any $\pi \in \Pi_i$, we solve the following

¹⁶Technically, this assumes that every agent is interacting in every element of the partition in every period, but the interaction is simply weighted by \mathbf{q} . Hence, there is a positive probability that i receives a link from every agent in i 's social group in every period despite the fact that they may not actually be interacting in that social group in that period. Furthermore, to derive equation (5.3) we implicitly assume that agents in π have the same degree as i at t .

difference equation

$$\begin{aligned} R^\pi(t+1) &= R^\pi(t) - [2q^\pi + (1-p^\pi)R^\pi(t) - (1-p^\pi)2q^\pi] \\ R^\pi(t+1) &= p^\pi [R^\pi(t) - 2q^\pi] \end{aligned} \quad (5.5)$$

where $p^\pi = \frac{1}{|\pi|} \sum_{i \in \pi} p_i$. The interpretation of equation (5.5) is straightforward. The number of remaining active agents in π at $t+1$ is simply the number of active agents in π at t less the number of agents that have either become idle or were linked with i . This includes the agents who were linked with i at t ($2q^\pi$) and those who have become idle at t ($(1-p^\pi)R^\pi(t)$) and excludes the ones who were linked with i at t and have become idle at t ($(1-p^\pi)2q^\pi$). For any agent i and any $\pi \in \Pi_i$, we can solve this with $R^\pi(0) = |\pi|$ to get

$$R^\pi(t) = |\pi|(p^\pi)^t + \frac{2q^\pi p^\pi ((p^\pi)^t - 1)}{1 - p^\pi} \quad (5.6)$$

Solving equation (5.6) for $R^\pi(T^\pi) = 0$, gives us the expected number of periods it takes i to form links with every agent in $\pi \in \Pi_i$, namely

$$T^\pi = \begin{cases} \frac{\ln\left(\frac{2q^\pi p^\pi}{2q^\pi p^\pi + (1-p^\pi)|\pi|}\right)}{\ln(p^\pi)} & \text{if } q^\pi > 0 \\ 0 & \text{otherwise} \end{cases} \quad (5.7)$$

The degree of agent i in period t is therefore given by equation (5.8) below

$$\begin{aligned} d_i(t) &= \sum_{\pi \in \Pi_i} d_i^\pi(t) \\ &= 2 \sum_{\pi \in \Pi_i} q^\pi [t \mathbf{1}(t \leq T^\pi) + T^\pi \mathbf{1}(t > T^\pi)] \end{aligned} \quad (5.8)$$

Note that $d_i(t)$ is a concave, piecewise linear function that is strictly increasing in the range $[0, \max_{\pi \in \Pi_i} \{T^\pi\}]$. This means that in our model, an active agent makes friends

at a decreasing rate over time. Given that preferences do not enter into our model, this feature is not exhibited because agents have decreasing marginal utility of friendships (Currarini et al., 2009), but rather because elements of the partitions Π_i , for every i , are gradually exhausted over time.

Since $d_i(t)$ is increasing, we can find its inverse in the range $[0, d_i(\max_{\pi \in \Pi_i} \{T^\pi\})]$, which is given by

$$d_i^{-1}(d) = t_i(d) = \frac{d - 2 \sum_{\pi \in \Pi_i} q^\pi T^\pi \mathbf{1}(d > d_i(T^\pi))}{2 \sum_{\pi \in \Pi_i} q^\pi \mathbf{1}(d \leq d_i(T^\pi))} \quad (5.9)$$

We now obtain $G_i(d)$ – the probability that agent i has degree at most d (degree distribution of agent i).

$$\Pr(d_i(t) \leq d) = \Pr(d_i^{-1}(d_i(t)) \leq d_i^{-1}(d)) = \Pr(t \leq t_i(d)) = G_i(d) \quad (5.10)$$

Since an agent i remains active exactly x periods with probability $p_i^x(1 - p_i)$, we have that

$$\Pr(t \leq x) = \sum_{t=0}^{t=x} p_i^t(1 - p_i) = 1 - p_i^{x+1} \quad (5.11)$$

Therefore, the degree distribution of agent i is given by

$$G_i(d) = \Pr(t \leq t_i(d)) = 1 - p_i^{t_i(d)+1} \quad (5.12)$$

Finally, the overall degree distribution $G(d)$ is the average of the degree distributions across all agents and is given by

$$G(d) = \frac{1}{|N|} \sum_{i \in N} \left(1 - p_i^{t_i(d)+1}\right) \quad (5.13)$$

Note that the overall degree distribution is approximately exponential.¹⁷ We discuss the implications of this in section 5.9.3. Henceforth, in order to keep the model parsimonious and reduce the number of parameters when it comes to calibrating it to the data, we shall assume that $p_i = p$ for all i .

The results derived above allow us to obtain the relationship between the size of the social group and degree. The following proposition shows that agents in larger social groups, *ceteris paribus*, have a higher degree.

Proposition 5.1. *Consider agents i and j such that $|\Gamma_i^r| = |\Gamma_j^r|$ for all $r \in S \setminus \{r'\}$, but $|\Gamma_i^{r'}| > |\Gamma_j^{r'}|$. Then, for all t , $d_i(t) \geq d_j(t)$ and for some t , the inequality is strict.*

Proof. Using equations (5.2) and (5.7), we obtain

$$T^{\pi_i(S)} = \frac{\ln \left(\frac{2p \left[\sum_{r \in S \cup \{0\}} \frac{q^r}{|\Gamma_i^r|} \right]}{2p \left[\sum_{r \in S \cup \{0\}} \frac{q^r}{|\Gamma_i^r|} \right] + (1-p)} \right)}{\ln(p)} \quad (5.14)$$

$T^{\pi_i(S)}$ is increasing in $|\Gamma_i^r|$ for every S such that $r \in S$. Therefore $T^{\pi_i(S)} = T^{\pi_j(S)}$ for every $S \setminus \{r'\}$ and $T^{\pi_i(S')} > T^{\pi_j(S')}$ for any S' such that $r' \in S'$. Recall that $d_i(t)$ is a concave, piecewise linear function that is strictly increasing in the range $[0, \max_{\pi \in \Pi_i} \{T^\pi\}]$. Hence, we have that $d_i(T^{\pi_i(S')}) > d_j(T^{\pi_j(S')})$ and the result follows. \square

5.4.2 Assortativity

The simple structure of the model allows us to derive a further property of the resulting social network.

Proposition 5.2. *The average degree of agent i 's friends is increasing in agent i 's degree.*

¹⁷The degree distribution of agent i is geometric. The exponential distribution is a continuous analogue of the geometric distribution.

Proof. Let $(d_i)_{i \in N}$ and $(d'_i)_{i \in N}$ be vectors of realised degrees in the network formation process and \mathcal{N}_i is the set of i 's friends. It suffices to show that $d'_i > d_i$ implies that $\sum_{j \in \mathcal{N}_i} d'_j > \sum_{j \in \mathcal{N}_i} d_j$. Since $d'_i > d_i$, we have that $t_i(d'_i) > t_i(d_i)$ (see Figure 5.3). Hence, for some $\pi \in \Pi_i$, it must be true that $d_i^\pi(t_i(d'_i)) - d_i^\pi(t_i(d_i)) > 0$, which also holds for any $j \in \pi$. Therefore, the change in total degree of i , given by $\sum_{\pi \in \Pi_i} [d_i^\pi(t_i(d'_i)) - d_i^\pi(t_i(d_i))] > 0$. Hence, for all $j \in \mathcal{N}_i$, $\sum_{\pi \in \Pi_j} [d_j^\pi(t_i(d'_i)) - d_j^\pi(t_i(d_i))] > 0$, which completes the proof. \square

5.4.3 Homophily

Homophily captures the tendency of agents to form links with those similar to themselves. We now present definitions for a class of homophily measures and show the relationship between them (McPherson et al., 2001). We then express homophily within the context of our model and derive several results that describe the dynamics of homophily in the link formation process.

5.4.3.1 Individual homophily

For any agent i , the *individual homophily index* in social category $r \in \mathcal{R}$ is given by

$$H_i^r = \frac{\text{number of friends of } i \text{ that share } k_i^r}{\text{number of friends of } i} \quad (5.15)$$

Let $W_k^r = \{j \in N | k_j^r = k\}$ be the set of all agents that have characteristic $k \in K^r$. We say that an agent exhibits no individual homophily in social category r if the individual homophily index equals to the fraction of agents in the population who have characteristic $k = k_i^r$ i.e. if $H_i^r = \frac{|W_k^r|}{|N|}$.

5.4.3.2 Group homophily

We now present a definition of group homophily which corresponds to Definition 1 in Currarini et al. (2009). For any characteristic k in social category r , the *group*

homophily index is given by

$$\mathbf{H}_k^r = \frac{\sum_{i \in W_k^r} \text{number of friends of } i \text{ that share } k_i^r}{\sum_{i \in W_k^r} \text{number of friends of } i} \quad (5.16)$$

which is the fraction of the total number of friendships that agents with characteristic k have made with agents who also have characteristic k . We say that a group exhibits homophilious behaviour in social category r if the group homophily index exceeds the fraction of agents in the population who have characteristic k i.e. if $\mathbf{H}_k^r > \frac{|W_k^r|}{|N|}$. Heterophilious behaviour is defined analogously as $\mathbf{H}_k^r < \frac{|W_k^r|}{|N|}$ (Definition 5 in Currarini et al. (2009)).¹⁸

It is easy to verify the following relationship between individual and group homophily

$$\sum_{i \in W_k^r} H_i^r \times \left[\frac{\text{number of friends of } i}{\sum_{i \in W_k^r} \text{number of friends of } i} \right] = \mathbf{H}_k^r \quad (5.17)$$

5.4.3.3 Homophily in our model

Let us define $\Pi_i^r = \{\pi_i(S) \in \Pi_i | r \in S\}$. This is the set of partition elements that contain agents who share i 's characteristic in social category r . Using equation (5.15),

¹⁸In a similar vein, we can define *inbreeding homophily* (Currarini et al., 2009). First, for any agent i , the *individual inbreeding homophily index* (IH_i^r) in social category $r \in \mathcal{R}$ is given by

$$IH_i^r = \frac{H_i^r - \frac{|W_k^r|}{|N|}}{1 - \frac{|W_k^r|}{|N|}}$$

which captures how homophilious an agent i is relative to how homophilious i could be given the number of agents who have characteristic k_i^r in the population. Now, for any characteristic k in social category r , the *group inbreeding homophily index* (Definition 6 in Currarini et al. (2009)) is given by

$$\mathbf{IH}_k^r = \frac{\mathbf{H}_k^r - \frac{|W_k^r|}{|N|}}{1 - \frac{|W_k^r|}{|N|}}$$

which captures how homophilious a group of agents is relative to how homophilious the group could be. We do not focus on this class of homophily measures in this chapter.

the individual homophily index in social category r of agent i in period t is

$$H_i^r(t) = \frac{\sum_{\pi \in \Pi_i^r} d_i^\pi(t)}{\sum_{\pi \in \Pi_i} d_i^\pi(t)} = \frac{\sum_{\pi \in \Pi_i^r} d_i^\pi(t)}{d_i(t)} \quad (5.18)$$

Following equations (5.16) and (5.17), we obtain the group homophily index for characteristic k in social category r in period t

$$\begin{aligned} \sum_{i \in W_k^r} H_i^r(t) \left[\frac{d_i(t)}{\sum_{i \in W_k^r} d_i(t)} \right] &= \\ \sum_{i \in W_k^r} \left[\frac{\sum_{\pi \in \Pi_i^r} d_i^\pi(t)}{d_i(t)} \right] \left[\frac{d_i(t)}{\sum_{i \in W_k^r} d_i(t)} \right] &= \\ \sum_{i \in W_k^r} \left[\frac{\sum_{\pi \in \Pi_i^r} d_i^\pi(t)}{\sum_{i \in W_k^r} d_i(t)} \right] &= \mathbf{H}_k^r(t) \end{aligned} \quad (5.19)$$

Finally, it will be useful to define a composition function $h_i^r(d) \equiv (H_i^r \circ t_i)(d)$ which expresses individual homophily as a function of degree rather than as a function of time.

5.4.3.4 Dynamics of homophily

We now explore the properties of $H_i^r(t)$. Let $T_i^L = \min_{\pi \in \Pi_i} \{T^\pi\}$, $T_i^M = \max_{\pi \in \Pi_i^r} \{T^\pi\}$, and $T_i^H = \max_{\pi \in \Pi_i} \{T^\pi\}$. Note that $T_i^L \leq T_i^M \leq T_i^H$ for all i .

Proposition 5.3. *The function $H_i^r(t)$ has the following form: (i) for $t \in (0, T_i^L)$, $H_i^r(t)$ is a constant; (ii) for $t \in [T_i^L, T_i^M)$, the slope of $H_i^r(t)$ is ambiguous; (iii) for $t \in [T_i^M, T_i^H)$, $H_i^r(t)$ is decreasing; (iv) for $t \in [T_i^H, \infty)$, $H_i^r(t)$ is a constant.*

Proof. First of all, note that both the numerator and the denominator are concave, non-decreasing, piecewise linear functions, and for any given t , the slope of the numerator is always less than that of the denominator. (i) For $t \in (0, T_i^L)$, both the numerator and the denominator are linear functions starting at the origin, with the denominator

having a steeper slope than the numerator. Hence, $H_i^r(t)$ is a constant. (ii) At T_i^L , there is a kink either (a) in the denominator alone or (b) both in the numerator and the denominator. In case (a), $H_i^r(t)$ would increase since the slope of the denominator falls, but in case (b), it is ambiguous (it is easy to find an example where $H_i^r(t)$ increases before decreasing again in this range). This reasoning applies every time there is such a kink, which occurs at every expected stopping time in the interval $[T_i^L, T_i^M]$. (iii) At T_M the numerator becomes flat, while the denominator is still increasing. This implies that $H_i^r(t)$ is decreasing in the interval $[T_i^M, T_i^H]$. (iv) Finally, at T_i^H , the denominator also becomes flat, which means that for every $t \geq T_i^H$, $H_i^r(t)$ is simply a constant divided by another constant. \square

Remark 1. $h_i^r(d) \equiv (H_i^r \circ t_i)(d)$ has a similar shape to $H_i^r(t)$. For this, it suffices to note that $t_i(d)$ is an increasing, piecewise linear function.


Figure 5.3 illustrates the general relationships between $H_i^r(t)$, $h_i^r(d)$, and $d_i(t)$ or $t_i(d)$, depending on whether we want to take degree or time as the exogenous variable. In the full model, the complex correlations between agents' social groups introduce non-monotonicities in the individual homophily indices. This means that a tighter characterisation of homophily patterns cannot be given in the general case. The value of $H_i^r(1)$ could, for example, be below the limiting value of the individual homophily index. However, in section 5.11.6 of the Appendix, we show that this is not the case for most parameter values in our data. Hence, homophily patterns presented in Figure 5.3 are fairly representative of the homophily patterns in our data.

In order to give a feel for the dynamics of homophily in our model, we can provide tight analytical results for the case where $q^r > 0$ for only one $r \in \mathcal{R}_+$.

Corollary 5.1. *Consider the case where $q^r > 0$ for only one $r \in \mathcal{R}_+$. Then*

- (1) $H_i^r(t)$ and $h_i^r(d)$ are (weakly) decreasing in their respective arguments.

Figure 5.3: Relationship between degree and individual homophily indices



(2) For a given t , individual homophily for agent i in social category r (weakly) decreases in $|\Gamma_i^r|$.

(3) The variance of $\mathbf{h}^r(d) = (h_i^r(d))_{i \in N}$ converges to a constant.

Proof. (1) Since $q^r > 0$ for only one $r \in \mathcal{R}_+$, one can verify that (i) $T^{\pi_i(S)} \in \{T^{\pi_i(\{r\})}, 0\}$ for any S such that $r \in S$, and (ii) $T^{\pi_i(S)} \in \{T^{\pi_i(\{0\})}, 0\}$ for any S such that $r \notin S$. Note that if $S' \subseteq S$, then $\sum_{r \in S'} \frac{q^r}{|\Gamma_i^r|} \leq \sum_{r \in S} \frac{q^r}{|\Gamma_i^r|}$ and therefore from equation (5.14), we obtain $T^{\pi_i(S')} \geq T^{\pi_i(S)}$ (since $\ln(p) < 0$). Hence, $\max_{\pi \in \Pi_i} \{T^\pi\} = T^{\pi_i(\{0\})}$, $\max_{\pi \in \Pi_i^r} \{T^\pi\} = T^{\pi_i(\{r\})}$ and $\min_{\pi \in \Pi_i} \{T^\pi\} \in \{T^{\pi_i(\{r\})}, 0\}$. From this, it follows that: either $T_i^L = T_i^M = T^{\pi_i(\{r\})}$ and $T_i^H = T^{\pi_i(\{0\})}$; or $T_i^L = 0$, $T_i^M = T^{\pi_i(\{r\})}$ and $T_i^H = T^{\pi_i(\{0\})}$. From Proposition 5.3, we have that $H_i^r(t)$ has the following form: (i) for $t \in [0, T_i^M)$, $H_i^r(t)$ is a constant; (ii) for $t \in [T_i^M, T_i^H)$, $H_i^r(t)$ is decreasing; (iii) for $t \in [T_i^H, \infty)$, $H_i^r(t)$ is a constant. The shape of $h_i^r(d)$ follows from this and from Remark 1.

(2) This follows immediately from Part (1) and Proposition 5.1.

(3) This follows immediately from Part (iv) of Proposition 5.3 and from Remark 1. □

5.5 Simulation results

We used the mean-field approximation method to derive the analytical expressions for the dynamics of the network formation process. As we mentioned in section 5.4, the accuracy of the mean-field approximation must first be tested against simulations. The simulation algorithm, which emulates the theoretical network formation process, is summarised in section 5.11.1 of the Appendix. We tested the analytical expressions for the degree distribution and the individual homophily index distribution against an average of 100 runs of the simulation for multiple parameter values. Our analytical degree distribution matches the simulated version exceptionally well. There is, however, some loss of accuracy at extreme values of the cumulative distribution of the individual homophily index. Nevertheless, we anticipated this in the theoretical model. Equation (5.18) makes it clear that the individual homophily index is unlikely to be 0 or 1. The individual homophily index is 0 when $\sum_{\pi \in \Pi_i^r} d_i^\pi(t) = 0$ i.e. only if $\Gamma_i^r = \emptyset$. This can only happen when an agent is alone in her social group. The individual homophily index could be 1 for the case of, say, gender in a women's college (i.e. $\sum_{\pi \in \Pi_i^r} d_i^\pi(t) = d_i(t)$). This is purely an artifact of the mean-field approximation of the individual degree. Despite these problems at the extremes, if the model is correct, we should expect a good prediction of the *average* of the individual homophily indices. Head-to-head plots and numerical results for both the degree distribution and homophily patterns are provided in section 5.8.2 below.¹⁹

¹⁹The plots in section 5.8.2 only show an example, but the tests of the analytical approximations against the simulations were run for a broad range of values.

5.6 Data

We use Facebook data first analysed by Traud et al. (2010) and Traud et al. (2012). The data represent a September 2005 cross-section of the complete structures of social connections on www.facebook.com *within* (but not across) the first ten American colleges and universities that joined Facebook. The raw data contain over 130,000 nodes (users) and over 5.6 million links (friendships). In order to join Facebook, each user had to have a valid college email address. We observe six social categories for each user: gender, year of graduation, major, minor, dorm, and high school. In order to protect personal identity, all the data were anonymised by Adam D'Angelo (former CTO of Facebook) and are represented by number codes.

Since all personal data were provided voluntarily, some users did not submit all their information. Testing our model requires us to observe major, minor, dorm, gender, and year of graduation for every user.²⁰ We dropped any user (and their links), who has not provided all the personal characteristics other than high school. In addition, some users were listed as faculty members and some students listed graduation years that were probably untruthful (e.g. 1926). We therefore dropped all faculty members and every user whose year of graduation is outside 2006-2009. Hence, in our data, we look only at students graduating between 2006 and 2009, who have supplied all the relevant personal characteristics (except high school).²¹

There are 27,454 users and 492,236 links in our cleaned dataset. The individual college names were provided in abbreviated form, however, we managed to back out

²⁰High school is also an interesting immutable social category, however, the relative group sizes within colleges are too small to allow for a meaningful analysis.

²¹Technically, this means we consider a non-random subsample of the data since there might be selection biases in data disclosure preferences. However, in section 5.11.5 of the Appendix, we show that the degree distributions in the cleaned datasets are very similar to the original datasets. Hence, we expect that our calibrated parameters should be close to the unbiased parameter estimates and that our comparative statics results should remain unchanged. Structural estimation of our model using the full data is a potential area of future research.

the names of all colleges using their tags from the order in which they appear in our dataset and the order in which they joined Facebook.²² The summary of the data is given in section 5.11.3 of the Appendix.

5.7 Tests and empirical observations

5.7.1 A representative college

Before we calibrate the model to the data, let us first get a feel for the general empirical patterns and the information contained in our dataset. Since there are ten separate networks, it is impractical to give the visual representations and detailed statistics for every college. Instead, whenever it is necessary, we focus on a representative college.²³ For example, Figure 5.4(a) shows the network for Harvard University (the first college to have Facebook) with nodes in the graph coloured by graduation year.²⁴ We can see that students from the same year group tend to cluster together. Another way of illustrating this would be by considering the adjacency matrix directly. In Figure 5.4(b), we plot the adjacency matrix with the students sorted by the year of graduation, where each point represents a link (Newman, 2010, p. 227). In section 5.11.6 of the Appendix, we also show that that the dynamics of homophily presented in Figure 5.3 hold quite generally.

5.7.2 All colleges

We also offer some tentative support for the dynamic predictions of our model. While the dataset is a cross section, we can look at the homophily degree patterns across year groups. This is clearly imperfect, but it provides some indication of whether the model will be able to match data in a panel dataset. Figure 5.5(a) shows that, on


²²Using, *inter alia*, a community edited public list: <http://www.quora.com/Facebook-Company-History/In-what-order-did-Facebook-open-to-college-and-university-campuses>

²³The Matlab and Python code and analogous results for any college are available upon request.


²⁴This was generated with Yifan Hu's Multilevel algorithm using the open-source *Gephi* graph visualisation software.

Figure 5.4: Harvard University Facebook network in September 2005

(a) Year of graduation: Red - 2009; Purple - 2008; Blue - 2007; Green - 2006



(b) Assortative matching by year of graduation



average, degree is non-decreasing across year groups (over time). Degree seems to fall for the students graduating in 2006, but the behavior of seniors may have differed slightly from the other cohorts since they were about to leave college when Facebook was introduced. Figure 5.5(c) shows that, as predicted by the model (see Corollary 5.1), (on average) the individual homophily index in year of graduation falls sharply as students enter later years. Figure 5.5(d) shows that gender homophily is roughly stable across the years, which is permitted by our model (see Proposition 5.3 and Corollary 5.1).²⁵ Figure 5.5(b) shows that more popular students are friends with other more popular students. This was the prediction of Proposition 5.2.

We also test Proposition 5.1, which states that, *ceteris paribus*, agents in larger social groups should have a higher degree. As a straightforward check of this proposition, for each college we regress the degree of each agent on the size of his class and dorm

$$d_i = \alpha + \beta_1 \times \text{class size}_i + \beta_2 \times \text{dorm size}_i + \epsilon \quad (5.20)$$

Table 5.1 reports the results for all ten colleges. We find that most coefficients are positive or not significantly different from zero in support of our model.

5.8 Model calibration


5.8.1 Empirical strategy

We calibrate our model against the data using the social categories identified in the Example. Using the available information in our data, we define agents i and j to be in the same class if and only if they have the same year of graduation and major or have the same year of graduation and minor. We assume that every agent i interacts in her class and dorm with respective probabilities q^1 and q^2 . The probability of interacting


²⁵Dashed line represent 99% Chebyshev confidence intervals.

Figure 5.5: Testing predictions of the model


(a) Degree across year groups



(b) Positive assortativity



(c) Year of graduation homophily across year groups



(d) Gender homophily across year groups




Table 5.1: Degree and class/dorm size

Dependent variable: agent's degree						
	class size	std. error	dorm size	std. error	const.	N
Harvard U.	0.151***	(0.023)	0.230***	(0.032)	9.243	1325
Columbia U.	0.095***	(0.019)	-0.018*	(0.010)	36.6	2663
Stanford U.	0.260***	(0.028)	0.010	(0.017)	36.23	2254
Yale U.	0.069**	(0.030)	0.065***	(0.021)	24.7	1431
Cornell U.	0.007	(0.015)	-0.004**	(0.002)	21.54	2509
Dartmouth	0.177***	(0.032)	-0.057*	(0.032)	37.09	1612
U. of Penn.	0.120***	(0.014)	-0.026***	(0.004)	38.51	3006
M.I.T.	0.043	(0.027)	-0.042***	(0.014)	43.23	1563
New York U.	0.059***	(0.009)	0.014***	(0.003)	25.79	5581
Boston U.	0.051***	(0.009)	0.002*	(0.001)	28.61	5510

Comment: Standard OLS regression with robust standard errors in parentheses

***/**/* denote rejection of $H_0 : \beta = 0$ at the 1/5/10% significant level respectively

with the gender and year social categories are set to zero ($q^3 = q^4 = 0$; see section 5.3.3 for a justification). Finally, $q^0 = 1 - q^1 - q^2$ is the proportion of time spent interacting with all other agents (their free time). Hence, the model has 4 parameters (namely q^0 , q^1 , q^2 , and p), but only 3 degrees of freedom.

In order to fit the model to the data (degree distribution and homophily patterns), we used a grid search on parameters q^0 , q^1 , q^2 , and p . For q^0 , q^1 , and q^2 , we took values from 0 to 1 in steps of 0.05. For p , we took values from 0.90 to 0.9975 in steps of 0.0025. For the degree distribution, we computed the analytical degree distribution, and, for homophily, we found the analytical homophily index of every agent i in gender and year of graduation as a function of i 's empirical degree at each point in the grid. Our goal is to fit the degree distribution and vectors of individual homophily indices as closely as possible to the actual data.

Since there may be a trade-off in fitting homophily patterns and degree distribution, we found best-fitting values \tilde{q}^0 , \tilde{q}^1 , \tilde{q}^2 , and \tilde{p} , which minimise an intuitive loss function

that measures the “overall error” in our model.²⁶ For each point $(\mathbf{q}, p) = (q^0, q^1, q^2, p)$ in the grid, we define the distance $\Delta^d(\mathbf{q}, p)$ between the empirical degree distribution $G(x; \mathbf{q}, p)$ and the analytical degree distribution $\hat{G}(x; \mathbf{q}, p)$ as

$$\Delta^d(\mathbf{q}, p) = \sum_{x=0}^{x=\max_{i \in N} \{d_i\}} \left(G(x; \mathbf{q}, p) - \hat{G}(x; \mathbf{q}, p) \right)^2 \quad (5.21)$$

and let $\Delta^d = (\Delta^d(\mathbf{q}, p))_{(\mathbf{q}, p)}$. Similarly, for each point in the grid, we define the distance between the empirical $(h_i^r(\mathbf{q}, p))_{i \in N}$ and the analytical $(\hat{h}_i^r(\mathbf{q}, p))_{i \in N}$ vectors of individual homophily indices as follows

$$\Delta^r(\mathbf{q}, p) = \sum_{i \in N} \left(h_i^r(\mathbf{q}, p) - \hat{h}_i^r(\mathbf{q}, p) \right)^2 \quad (5.22)$$

as well as $\Delta^r = (\Delta^r(\mathbf{q}, p))_{(\mathbf{q}, p)}$. We would like to minimise the following loss function with respect to (\mathbf{q}, p)

$$\mathcal{L}(\mathbf{q}, p) = \left[\frac{\Delta^d(\mathbf{q}, p)}{\|\Delta^d\|} \times \frac{\Delta^3(\mathbf{q}, p)}{\|\Delta^3\|} \times \frac{\Delta^4(\mathbf{q}, p)}{\|\Delta^4\|} \right] \quad (5.23)$$

where $\|\Delta\|$ is the Euclidean norm of Δ . The normalisations guarantee that the distances are comparable across the various components of the loss function. Furthermore, note that the loss function puts equal weight on the normalised distances between the empirical and analytical degree distribution and between vectors of the individual homophily indices.

We ranked the 8680 grid points (\mathbf{q}, p) starting with the one that minimises $\mathcal{L}(\mathbf{q}, p)$. Since the grid search is necessarily coarser than a full optimisation, we wanted to avoid

²⁶In principle, one could define any sensible loss function. We opted for a Cobb-Douglas functional form with equal weights on the arguments. We could have also used the Generalised Method of Moments (GMM). However, the vectors of individual homophily indices for any college are of length N , whereas the analytical cumulative degree distributions may be of a different length. Implementing GMM appropriately would require the moment vectors to be of equal length, but reducing the vectors to the same length coarsens data and worsens fit.

the possibility of finding the highest ranked point by chance. That is, an isolated point could have been picked as a global minimum simply because of the way in which the grid was overlaid on the loss function. We developed a simple method to try to pin down the global minimum more robustly. Our algorithm identified sets of points (among the top 100 of the possible 8680) that are ‘near’ each other in the grid. These sets were ranked according to value of loss function at the points within each set. We selected the best point within the highest ranking set. The algorithm always selected one of the top two points among the top 100 possible points. The method is outlined in section 5.11.2 of the Appendix.

5.8.2 Results

We are interested in the structural properties of Facebook networks, such as degree distribution and clustering, as well as homophily in gender and year group, and in testing how closely our model reproduces them. Technically, our model and the various definitions of homophily allow us to measure homophily in any social category. However, characteristics within certain categories could, in principle, be chosen endogenously by the agents. For example, students can change their major depending on what major their friends have chosen (see our discussion of endogenous characteristics in section 5.9.4). Since our model does not account for such a feedback mechanism within the characteristics, we only consider our homophily results for immutable social categories in our dataset. These happen to be gender and year of graduation.

We ran model simulations for every network at its best-fitting parameter values, which minimised its loss function $\mathcal{L}(\mathbf{q}, p)$ according to the robust grid search algorithm. The table in section 5.11.4 of the Appendix presents the results for the first ten colleges that joined Facebook. It shows that our model predicts average individual homophily and the average individual clustering coefficient (see Jackson (2008, p. 35) for a stan-

dard definition) very well.²⁷ Remarkably, the clustering results from our simulations fit the empirical results even though clustering does not enter into the loss function. A simple visual representation of the results table in section 5.11.4 is given in Figure 5.7.²⁸

Despite differences in the collegiate life of American universities, the best-fitting parameter values suggest that students spend a larger proportion of time interacting with students in their class than in their dorm. Nevertheless, there are observable heterogeneities in the best-fitting parameter values across the colleges, which indicates that the model is sufficiently flexible to accommodate for them. For example, M.I.T. students appear to spend more time making friends in their dorms relative to Harvard students. It is worth noting that recently Shaw et al. (2011) also obtained this qualitative result using different methods on the same dataset.

We also look at how the best-fitting parameter values change across year groups. Figure 5.6 shows that as students go through college, less time is allocated to making friends in class and more to dorm. This is intuitive: most freshmen are allocated dorms (randomly), while many seniors self-select into dorms with their friends.

The table in section 5.11.4 reports results on average statistics and by itself provides no indication of how well our model fits the *full distributions* of degree, individual homophily indices, and individual clustering coefficients. For this we need to look at individual representative colleges. Plots in Figure 5.8 show the empirical, analytical, and simulated degree, individual homophily, and individual clustering distributions for Harvard University. The figures make it clear that our model does not fit only the average statistics, but also entire distributions surprisingly well. Furthermore, the fits are representative of the analogous plots for the other colleges.

²⁷The simulated values are taken as an average over 100 runs of the model.

²⁸In order to avoid making any assumptions about the distributions, we estimated standard errors around the empirical averages non-parametrically. Figure 5.7 therefore represents the Chebyshev confidence intervals at the 95% and 99% levels.

Figure 5.6: Best-fitting parameter values by year-group


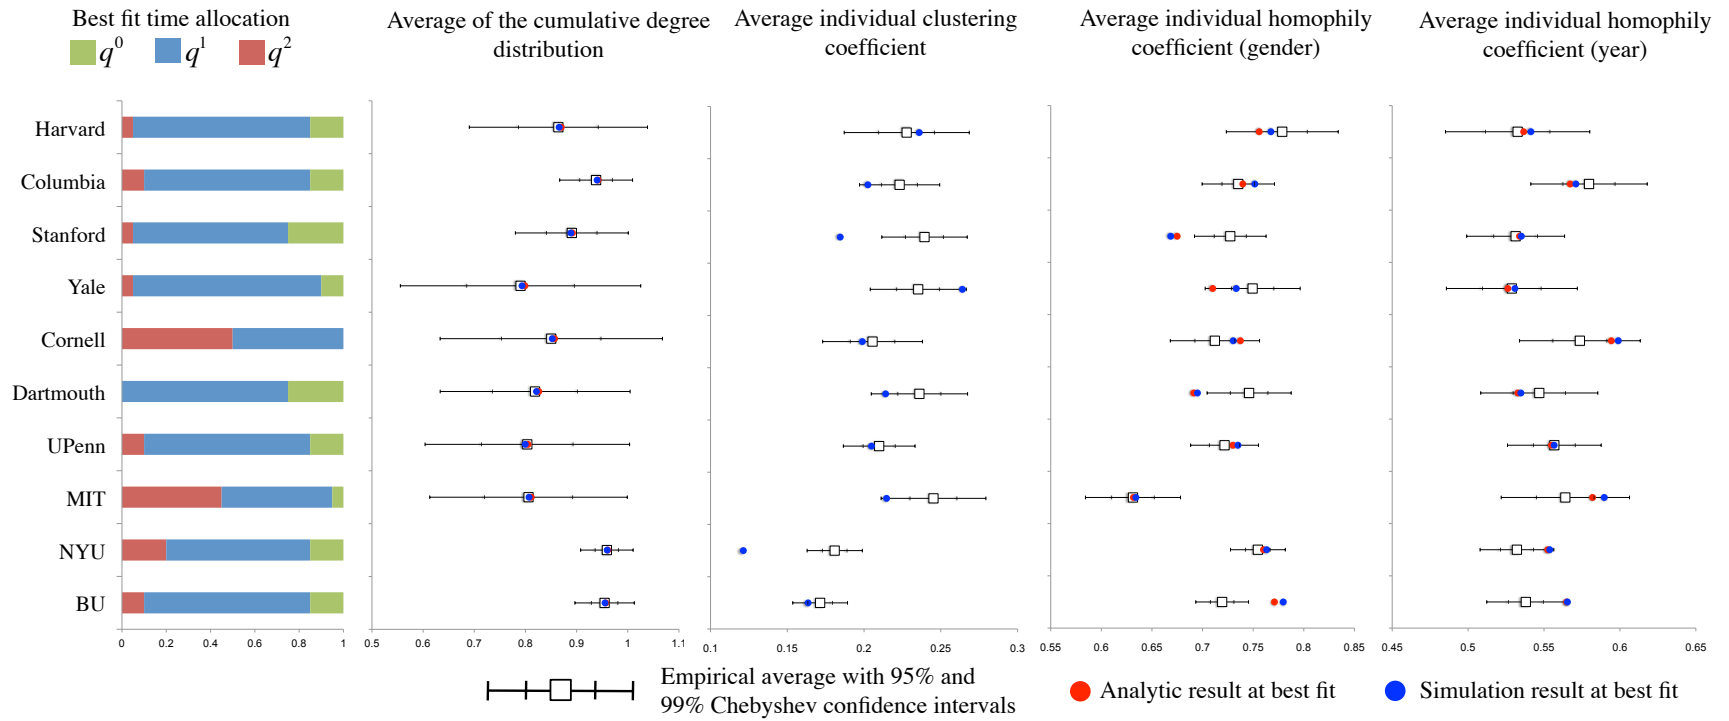


Figure 5.7: Illustration of results



5.9 Discussion

Our model lends itself to several potential extensions. In richer and more complex network formation processes, some of the extensions proposed below may be useful to obtain more accurate results on the network properties. As we mentioned above, the results of the model depend crucially on the expected stopping times, which are determined by equation (5.5). It should therefore be unsurprising that various extensions to the model involve modifying this equation.

5.9.1 Arrival of new nodes

So far we have ignored the arrival of new agents into the social network as we chose to give the simplest possible exposition of our model. However, incorporating this feature is straightforward. Suppose that the network formation process remains exactly the same as before but a new agent arrives in every period. Let us fix the distribution of characteristics of the population at $t = 0$. The characteristics of every new agent are always drawn randomly for every social category according to this initial distribution. In this case, for any existing agent i in the network, the probability that the new agent has characteristics of agents in $\pi_i(S)$ for any $S \in \mathcal{R}$ is

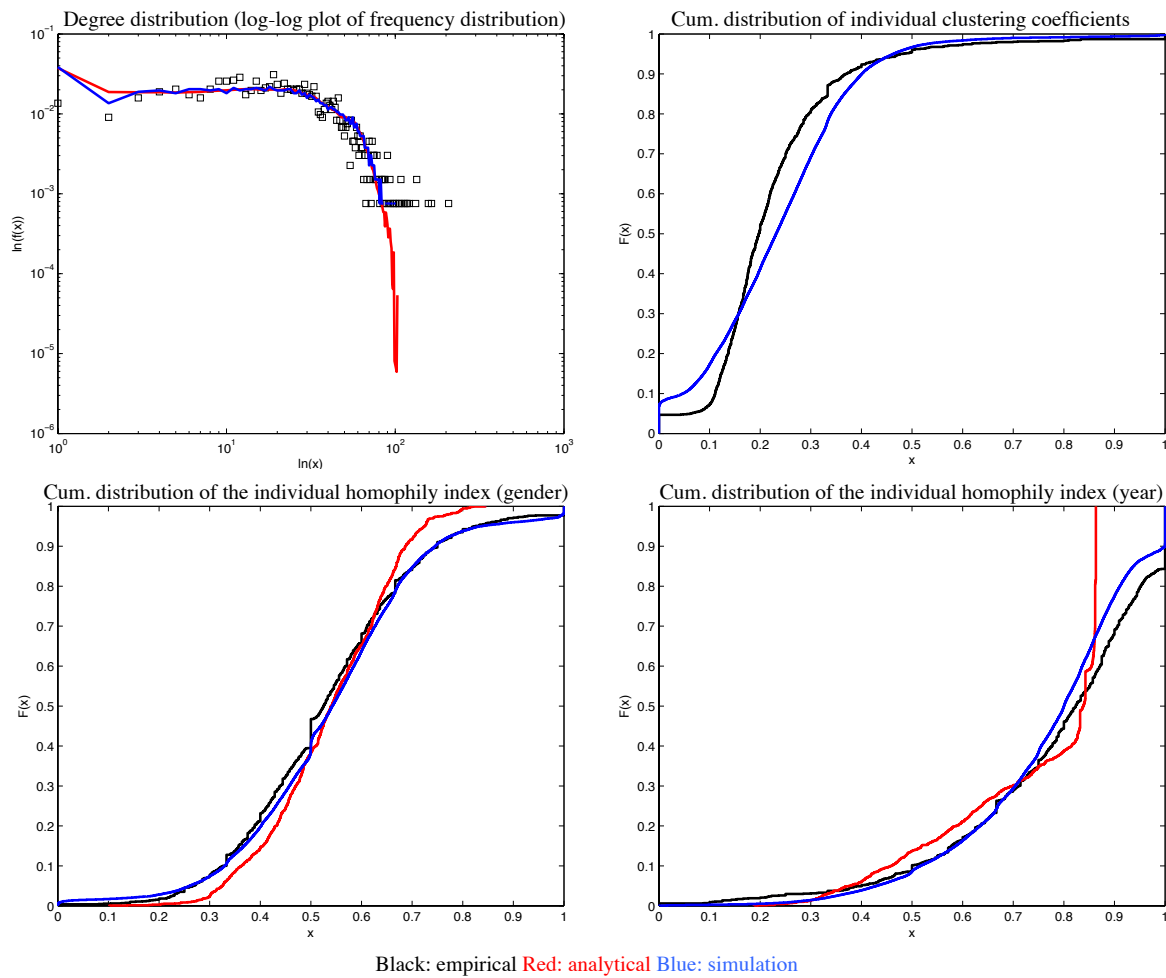
$$\mathcal{P}^{\pi_i(S)} = \prod_{r \in S} \frac{|\Gamma_i^r|}{|N|} \times \prod_{r \in \mathcal{R} \setminus S} \left(1 - \frac{|\Gamma_i^r|}{|N|}\right) \quad (5.24)$$

Using this fixed probability, equation (5.5) becomes

$$R^\pi(t+1) = R^\pi(t) + \mathcal{P}^\pi - [2q^\pi + (1-p)(R^\pi(t) + \mathcal{P}^\pi) - (1-p)2q^\pi] \quad (5.25)$$

$$R^\pi(t+1) = p[R^\pi(t) + \mathcal{P}^\pi - 2q^\pi] \quad (5.26)$$

Figure 5.8: Structural properties of the Facebook network at Harvard University



The intuition for this is that the remaining active agents in $\pi \in \Pi_i$ are the ones that remain active from: (i) those that were active in the previous period, as well as (ii) the new node (arriving with probability \mathcal{P}^π into this partition element), less (iii) the agents that i linked to in the previous period. Solving this once again for $R^\pi(0) = |\pi|$ yields

$$R^\pi(t) = |\pi|p^t + \frac{(2q^\pi - \mathcal{P}^\pi)p(p^t - 1)}{1 - p} \quad (5.27)$$

Finally, we can obtain the expected stopping time in π by solving $R^\pi(t) = 0$

$$T^\pi = \begin{cases} \frac{\ln\left(\frac{(2q^\pi - \mathcal{P}^\pi)p}{(2q^\pi - \mathcal{P}^\pi)p + (1-p)|\pi|}\right)}{\ln(p)} & \text{if } 2q^\pi - \mathcal{P}^\pi > 0 \\ 0 & \text{otherwise} \end{cases} \quad (5.28)$$

It is worth observing that the system is well defined only if the new nodes arrive at a slow enough rate ($\mathcal{P}^\pi < 2q^\pi$). If the new nodes arrive too quickly, stopping times become infinite. The rest of model can be solved using methods from section 5.4.

5.9.2 Endogenous probability of idleness

In the model, we have assumed that the probability $1 - p_i$ of an agent becoming idle in any given period is constant. It is reasonable to suppose that this probability depends on time, so instead we have $p_i(t)$. This would modify equations (5.5) and (5.11) respectively as follows

$$R^\pi(t+1) = p^\pi(t) [R^\pi(t) - 2q^\pi] \quad (5.29)$$

$$\Pr(t \leq x) = \sum_{t=0}^{t=x} \prod_{z=0}^{z=x} p_i(z) [1 - p_i(x+1)] \quad (5.30)$$

Although this introduces some difficulties for deriving the analytical expressions for the degree distribution for a general $p_i(t)$, these expressions will easily generate the appropriate numerical solutions.

5.9.3 Preferential attachment

Since the degree distribution in many networks follows a power law, Price (1976), Barabási and Albert (1999), and Jackson and Rogers (2007) suggested introducing preferential attachment into the network formation process in order to reproduce this property. This means that nodes in a network link to each other with a probability that is proportional to their degree.

We find that our model with *uniform random attachment* performs well against the data. In fact, Jackson (2008, p. 65) observes that

some of the more purely social networks have parameters that indicate much higher levels of random link formation, which are very far from satisfying a power law. In fact, the degree distribution of the romance network among high school students is essentially the same as that of a purely random network.

It is nevertheless possible to introduce preferential attachment into our model. Equation (5.3) becomes

$$\Delta d_i^\pi(t) = q^\pi \left(1 + d_i^\pi(t) \frac{R^\pi(t)}{R^\pi(t)} \right) \mathbf{1}(t \leq T^\pi) = (1 + d_i^\pi(t)) q^\pi \times \mathbf{1}(t \leq T^\pi) \quad (5.31)$$

Even though agent i 's out-link is made according to preferential attachment, it does not matter to whom it is made. However, the in-link is no longer made uniformly with probability $\frac{1}{R^\pi(t)}$, but instead with probability $\frac{d_i^\pi(t)}{R^\pi(t)}$, which is proportional to i 's degree. As before solving with $d_i^\pi(0) = 0$ yields

$$d_i^\pi(t) = (1 + q^\pi)^{[t\mathbf{1}(t \leq T^\pi) + T^\pi \mathbf{1}(t > T^\pi)]} - 1 \quad (5.32)$$

Now the analogue of equation (5.5) is

$$R^\pi(t+1) = p [R^\pi(t) - (1 + q^\pi)^t + 1] \quad (5.33)$$

assuming for simplicity that $p_i = p$ for all i . Setting $R^\pi(0) = |\pi|$ produces a rather unwieldy result

$$\begin{aligned}
R^\pi(t) = & \frac{p^{t+2}|\pi| - p^{t+1}[(2 + q^\pi)|\pi| + q^\pi] + p^t(1 + q^\pi)|\pi|}{(p - 1)(p - q^\pi - 1)} + \\
& + \frac{p^2[(1 + q^\pi)^t - 1] + p[1 + q^\pi - (1 + q^\pi)^t]}{(p - 1)(p - q^\pi - 1)} \tag{5.34}
\end{aligned}$$

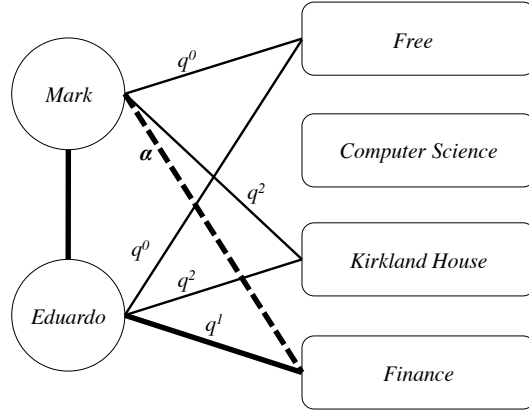
Solving for $R^\pi(T^\pi) = 0$, is possible numerically and the rest of derivation, once again, follows the steps outlined in section 5.4.

5.9.4 Endogenous characteristics

Perhaps the most interesting extension of the model is to consider what happens when \mathbf{k}_i is made endogenous (Bramoullé et al. (2012) and Boucher (2012) also make this point). Let us, once again, think about the model as an affiliation network discussed in section 5.3.5. New links can form in ways other than by focal closure. If an agent i is linked to another agent j who has a particular membership $k \in K^r$, then there is a positive probability that a link will form between i and $k \in K^r$. Figure 5.9 shows that Mark wants to join the Finance membership because his friend Eduardo is already a member. This is called *membership closure*.

The endogenous determination of characteristics in our model would be neatly captured by membership closure with a twist. In a standard affiliation model, Mark would create a new link to the Finance membership *in addition to* his link to Computer Science, whereas in our set-up, Mark would first delete his Computer Science link. The remaining conceptual difficulty would be to determine precisely what α – the probability with which Mark switches memberships – is. One possibility is to tie α to the membership of Mark’s friends: in each period, Mark may have a constant probability β^r of switching his memberships in some social category K^r , and the probability of choosing a new membership in K^r could be set in proportion to the number of Mark’s

Figure 5.9: Model with endogenous characteristics as an affiliation network



friends who have that membership.

5.10 Conclusion

We presented a dynamic network formation model, which provides rich microfoundations for the macroscopic properties of online social networks. Homophily patterns arise from random interaction within social groups. The analytical results of our parsimonious model find good support in the data. We were also able to estimate how much time agents spend in particular social groups. The model is flexible enough to allow for a variety of extensions. There is still scope for further theoretical work, including finding closed-form expressions to the clustering measures and diameter.

The model has some interesting implications for policy design. Suppose that the policy objective is to diffuse information about the quality of a particular product as quickly as possible and that agents learn by averaging signals about product quality from their neighbours (this is known as DeGroot learning). Golub and Jackson (2012) showed that homophily in a (random) network slows down the speed of DeGroot learning. In our model, agents who have been in the network the longest have the highest degree and are most heterophilous; therefore, information about the product would

travel fastest if the diffusion process began with these agents. Alternatively, it may indeed be effective to target the newest arrivals to the network because their homophily often increases as they make their first friendships.

5.11 Appendix

5.11.1 Simulation algorithm

Input:

$|N| \times R$ matrix \mathcal{M} where each row is the vector of characteristics \mathbf{k}_i .

Vector \mathbf{q} .

Initialise:

Empty adjacency matrix \mathbf{A} with elements a_{ij} .

Let L be the list of all agents.

Using \mathcal{M} , find $\{\Gamma_i^r | r \in \mathcal{R}\}$ for all $i \in N$.

1. while L is non-empty do
2. every agent in L becomes idle with probability p
3. L is now the list of remaining active agents in random order
4. for every i in L do
5. select an $r \in \mathcal{R}$ at random according to \mathbf{q}
6. if $Z = \Gamma_i^r \cap L \cap \{j \in N | a_{ij} = 0\} \neq \emptyset$ then
7. pick an agent j uniformly at random from Z
8. create edges ij and ji in \mathbf{A}
9. else continue to next agent in L
10. end if
11. end for
12. return \mathbf{A}

5.11.2 Algorithm for finding robust points in the grid search

Input:

Q where each row is a vector (\mathbf{q}, p) , and the rows are ordered by the value they induce in $\mathcal{L}(\mathbf{q}, p)$, from lowest at the top to highest at the bottom. \mathbf{Q} is the 100-by-4 matrix consisting of the top 100 row vectors of Q . $(\mathbf{q}, p)_k = (q_k^0, q_k^1, q_k^2, p_k) \in \mathbf{Q}$ denotes the k th row vector in \mathbf{Q} .

Initialise:

S is a 1-by-100 vector of scores, $\delta_q = 0.1$, $\delta_p = 0.05$.

1. for k from 1 to 100 do
2. $S(k) = |\{(\mathbf{q}, p)_j \in \mathbf{Q} | p_j \in [p_k - \delta_p, p_k + \delta_p] \ \& \ \forall i \in \{0, 1, 2\}, q_j^i \in [q_k^i - \delta_q, q_k^i + \delta_q]\}|$
3. end for
4. for k from 1 to 99 do
5. if $S(k) > S(k+1)$ then
6. break
7. else
8. continue
9. end if
10. end for
11. return $(\mathbf{q}, p)_k$

5.11.3 Data description

	Raw	Raw			Average			Ave.	Ave.	Ave.	Ave.
	number	number			degree	Women	Men	major	minor	dorm	class
College	of nodes	of edges	Nodes	Edges				size	size	size	size
Harvard U.	15,126	824,617	1,325	18,608	28.1	567	758	23.2	22.5	42.7	46.9
Columbia U.	11,770	444,333	2,663	52,697	39.6	1,573	1090	29.6	29.9	54.3	65.7
Stanford U.	11,621	568,330	2,254	55,124	48.9	1,043	1211	30.9	30.1	25.6	55.0
Yale U.	8,578	405,450	1,431	23,847	33.3	639	792	19.6	19.1	68.1	38.2
Cornell U.	18,660	790,777	2,509	26,653	21.2	1,078	1431	27.6	24.6	20.6	51.6
Dartmouth	7,694	304,076	1,612	34,030	42.2	780	832	29.9	29.3	23.0	45.0
U. of Penn.	14,916	686,501	3,006	60,516	40.3	1,417	1589	28.4	27.1	50.9	77.0
M.I.T.	6,440	251,252	1,563	32,751	41.9	626	937	44.7	37.2	26.1	58.2
New York U.	21,679	715,715	5,581	95,968	34.4	3,345	2236	53.7	52.2	101.5	99.7
Boston U.	19,700	637,528	5,510	92,042	33.4	3,355	2155	37.5	34.7	91.8	90.8
Average	13,618	562,858	2,745	49,224	36.3	1,442	1,303	32.5	30.7	50.5	62.8

5.11.4 Results

College	\tilde{q}^1	\tilde{q}^2	\tilde{q}^0	\tilde{p}	$\langle h_i^3 \rangle$	$\langle \hat{h}_i^3 \rangle$	$\langle \tilde{h}_i^3 \rangle$	$\langle h_i^4 \rangle$	$\langle \hat{h}_i^4 \rangle$	$\langle \tilde{h}_i^4 \rangle$	C	\hat{C}
Harvard U.	0.80	0.05	0.15	0.9625	0.53	0.54	0.54	0.78	0.76	0.77	0.23	0.24
Columbia U.	0.75	0.10	0.15	0.9700	0.58	0.57	0.57	0.74	0.74	0.75	0.22	0.20
Stanford U.	0.70	0.05	0.25	0.9775	0.53	0.53	0.54	0.73	0.67	0.67	0.24	0.18
Yale U.	0.85	0.05	0.10	0.9775	0.53	0.53	0.53	0.75	0.71	0.73	0.24	0.26
Cornell U.	0.50	0.50	0.00	0.9400	0.57	0.59	0.60	0.71	0.74	0.73	0.21	0.20
Dartmouth	0.75	0.00	0.25	0.9705	0.55	0.53	0.53	0.75	0.69	0.69	0.24	0.21
U. of Penn.	0.75	0.10	0.15	0.9725	0.56	0.55	0.56	0.72	0.73	0.73	0.21	0.20
M.I.T.	0.50	0.45	0.05	0.9700	0.56	0.58	0.59	0.63	0.63	0.63	0.25	0.21
New York U.	0.65	0.20	0.15	0.9550	0.53	0.55	0.55	0.75	0.76	0.76	0.18	0.12
Boston U.	0.75	0.10	0.15	0.9575	0.54	0.56	0.57	0.72	0.77	0.78	0.17	0.16

\tilde{q}^1 : best-fitting proportion of time spent in class

\tilde{q}^2 : best-fitting proportion of time spent in dorm

\tilde{q}^0 : best-fitting proportion of time spent as free time

\tilde{p} : best-fitting probability of remaining active in any given period

C : average of empirical individual clustering coefficients

\hat{C} : average of simulated individual clustering coefficients

$\langle h_i^3 \rangle$: average of empirical individual homophily indices for gender

$\langle \tilde{h}_i^3 \rangle$: average of analytical individual homophily indices for gender

$\langle \hat{h}_i^3 \rangle$: average of simulated individual homophily indices for gender

$\langle h_i^4 \rangle$: average of empirical individual homophily indices for graduation year

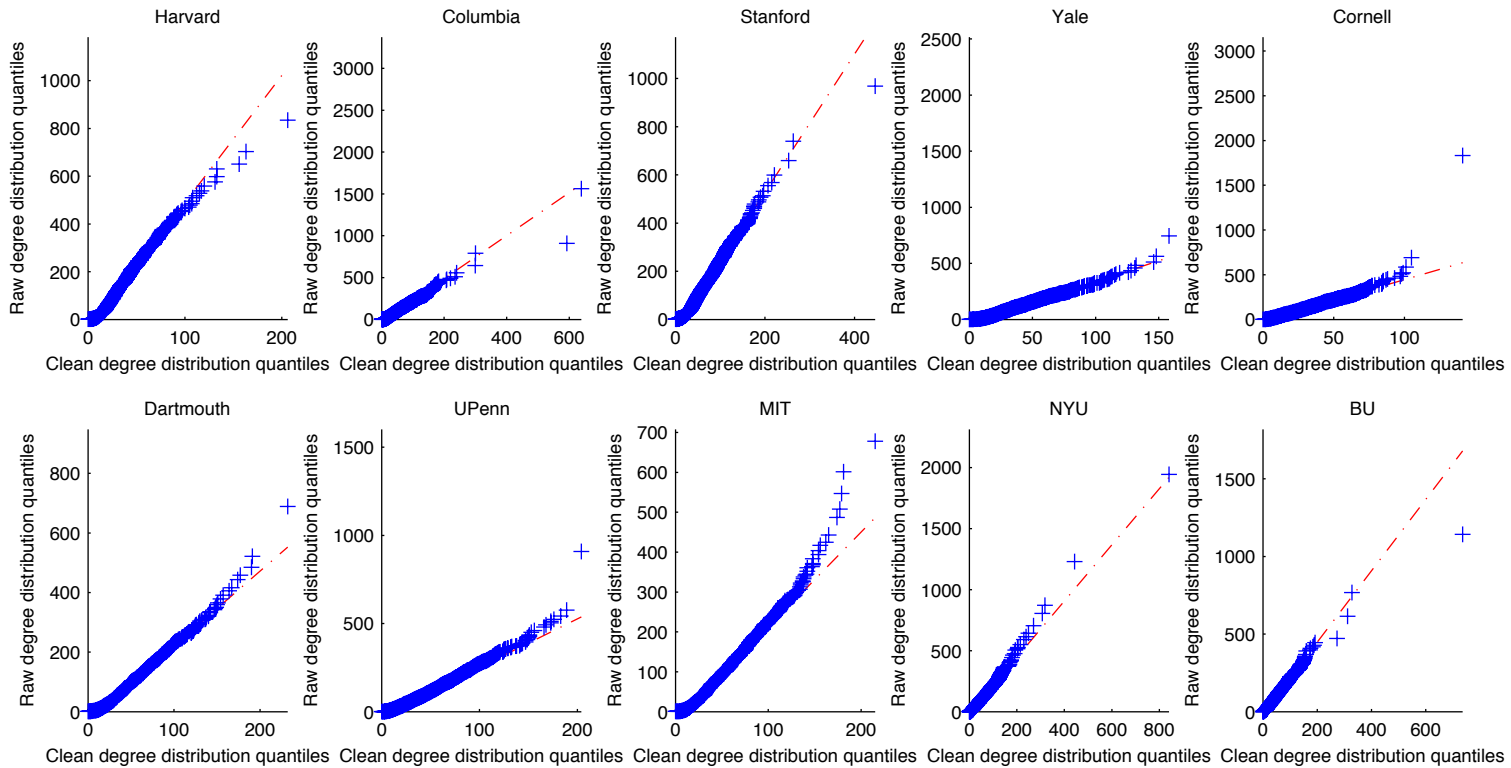
$\langle \tilde{h}_i^4 \rangle$: average of analytical individual homophily indices for graduation year

$\langle \hat{h}_i^4 \rangle$: average of simulated individual homophily indices for graduation year

5.11.5 Degree distributions in cleaned and raw data

This section presents the Q-Q plots for cleaned and raw datasets. A Q-Q plot shows the comparison between quantiles of the cleaned and raw degree distributions for a particular college. Two similar degree distributions should lie along the dashed-dotted $y = x$ line.

Figure 5.10: Degree distributions in cleaned and raw data



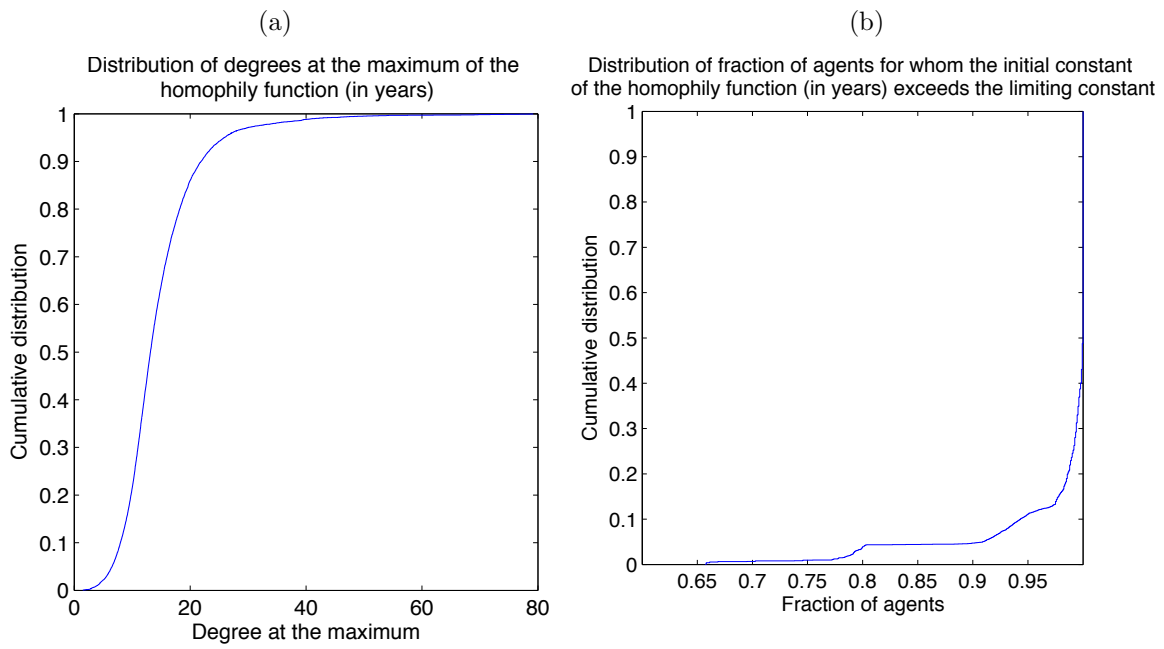
5.11.6 Dynamics of homophily across the grid space

In this section, we consider Harvard and individual homophily for year of graduation.

In Figure 5.11(a) for every grid point and every agent, we calculate what the degree of the agent is when his maximum level of individual homophily for graduation year is reached. We plot the cumulative distribution of this degree. We find that the median degree at which maximum individual homophily for graduation year is reached is 15 (average degree is 28). Therefore, individual homophily for year of graduation does not peak in the first period for most agents and parameter values.

In Figure 5.11(b) for every grid point, we calculate the proportion of agents for whom the initial constant level of the individual homophily for graduation year function exceeds the limiting constant (see Proposition 5.3). Here, for all grid points, the initial level of individual homophily for graduation year is higher than the limiting constant for at least a third of the agents. Additionally, for 90 percent of grid points, 95 percent of agents start out with a higher level of individual homophily for graduation year than the limiting constant. Results for any other college and for individual homophily for gender are comparable.

Figure 5.11: Dynamics of homophily



References

- Akerlof, G. A. and R. E. Kranton (2010). *Identity Economics: How Our Identities Shape Our Work, Wages, and Well-Being*. Princeton University Press.
- Banerjee, A., A. G. Chandrasekhar, E. Duflo, and M. O. Jackson (2012). The diffusion of microfinance. Working Paper 17743, NBER.
- Barabási, A.-L. and R. Albert (1999). Emergence of scaling in random networks. *Science* 286(5439), 509–512.
- Barabási, A.-L., R. Albert, and H. Jeong (1999). Mean-field theory for scale-free random networks. *Physica A* 272(1-2), 173–187.
- Bianconi, G., P. Pin, and M. Marsili (2009). Assessing the relevance of node features for network structure. *Proceedings of the National Academy of Sciences* 106(28), 11433–11438.
- Botha, L. and S. Kroon (2010). A community-based model of online social networks. In *4th SNA-KDD Workshop '10*.
- Boucher, V. (2012). Structural homophily. Working paper, Université de Montréal.
- Bramoullé, Y., S. Currarini, M. O. Jackson, P. Pin, and B. W. Rogers (2012). Homophily and long-run integration in social networks. *Journal of Economic Theory* 147(5), 1754–1786.
- Breiger, R. L. (1974). The duality of persons and groups. *Social Forces* 53(2), 181–190.
- Currarini, S., M. O. Jackson, and P. Pin (2009). An economic model of friendship: Homophily, minorities, and segregation. *Econometrica* 77(4), 1003–1045.
- Currarini, S., M. O. Jackson, and P. Pin (2010). Identifying the roles of race-based choice and chance in high school friendship network formation. *Proceedings of the National Academy of Sciences* 107(11), 4857–4861.
- de Marti, J. and Y. Zenou (2011, March). Identity and social distance in friendship formation. Mimeo.
- Easley, D. and J. Kleinberg (2010). *Networks, Crowds, and Markets: Reasoning about a highly connected world*. Cambridge University Press.
- Falk, A. and A. Ichino (2006). Clean evidence on peer effects. *Journal of Labor Economics* 24(1), 39–57.

- Feld, S. L. (1981). The focused organization of social ties. *American Journal of Sociology* 86(5), 1015–1035.
- Golub, B. and M. O. Jackson (2012). How homophily affects diffusion and learning in networks. *Quarterly Journal of Economics* 127(3), 1287–1338.
- Goyal, S. (2009). *Connections: An Introduction to the Economics of Networks*. Princeton University Press.
- Iijima, R. and Y. Kamada (2013, May). Social distance and network structures. Working paper.
- Jackson, M. O. (2008). *Social and Economic Networks*. Princeton University Press.
- Jackson, M. O. and B. W. Rogers (2007). Meeting strangers and friends of friends: How random are social networks? *American Economic Review* 70(3), 890–915.
- Kandel, D. B. (1978). Homophily, selection, and socialization in adolescent friendships. *American Journal of Sociology* 84(2), 427–436.
- Kremer, M. and D. Levy (2008). Peer effects and alcohol use among college students. *Journal of Economic Perspectives* 22(3), 189–206.
- Kumar, R., J. Novak, and A. Tomkins (2010). Structure and evolution of online social networks. In *Proceedings of the 11th ACM International Conference on Knowledge Discovery and Data Mining*, pp. 611–617.
- Lattanzi, S. and D. Sivakumar (2009). Affiliation networks. In *STOC' 09 Proceedings of the 41st Annual ACM Symposium on Theory of Computing*.
- Leskovec, J., J. Kleinberg, and C. Faloutsos (2005). Graphs over time: Densification laws, shrinking diameters and possible explanations. In *KDD '05 Proceedings of the Eleventh ACM SIGKDD International Conference on Knowledge Discovery in Data Mining*.
- Leskovec, J., K. J. Lang, A. Dasgupta, and M. W. Mahoney (2008). Statistical properties of community structure in large social and information networks. In *WWW '08 Proceedings of the 17th International Conference on World Wide Web*.
- Lewis, K., M. Gonzalez, and J. Kaufman (2012). Social selection and peer influence in an online social network. *Proceedings of the National Academy of Sciences* 109(1), 68–72.
- Lewis, K., J. Kaufman, M. Gonzalez, A. Wimmer, and N. Christakis (2008). Tastes, ties, and time: A new social network dataset using Facebook.com. *Social Networks* 30(4), 330–342.

- Mayer, A. and S. L. Puller (2008). The old boy (and girl) network: Social network formation on university campuses. *Journal of Public Economics* 92(1-2), 329–347.
- McPherson, M., L. Smith-Lovin, and J. M. Cook (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology* 27, 415–444.
- Moody, J. (2001). Race, school integration, and friendship segregation in America. *American Journal of Sociology* 107(3), 679–716.
- Mouw, T. and B. Entwisle (2006). Residential segregation and interracial friendship in schools. *American Journal of Sociology* 112(2), 394–441.
- Newman, M. E. J. (2010). *Networks: An Introduction*. Oxford University Press: Oxford.
- Newman, M. E. J., D. J. Watts, and S. H. Strogatz (2002). Random graph models of social networks. *Proceedings of the National Academy of Sciences* 99(Supplement 1), 2566–2572.
- Pempek, T. A., Y. A. Yermolayeva, and S. L. Calvert (2009). College students’ social networking experiences on Facebook. *Journal of Applied Developmental Psychology* 30(3), 227–238.
- Price, D. D. S. (1976). A general theory of bibliometric and other cumulative advantage processes. *Journal of the American Society for Information Science* 27(5), 292–306.
- Sacerdote, B. (2001). Peer effects with random assignment: Results for Dartmouth roommates. *Quarterly Journal of Economics* 116(2), 681–704.
- Sen, A. (2006). *Identity and Violence: Illusion of Destiny*. Allen Lane.
- Shaw, B., B. Huang, and T. Jebara (2011). Learning a distance metric from a network. In *Neural Information Processing Systems*.
- Shrum, W., N. H. Cheek Jr., and S. M. Hunter (1988). Friendship in school: Gender and racial homophily. *Sociology of Education* 61(4), 227–239.
- Traud, A. L., E. D. Kelsic, P. J. Mucha, and M. A. Porter (2010). Comparing community structure to characteristics in online collegiate social networks. *SIAM Review* 53(3), 526–543.
- Traud, A. L., P. J. Mucha, and M. A. Porter (2012). Social structure of Facebook networks. *Physica A* 391(16), 4165–4180.
- Wimmer, A. and K. Lewis (2010). Beyond and below racial homophily: ERG models of a friendship network documented on Facebook. *American Journal of Sociology* 116(2), 583–642.

Xiang, R., J. Neville, and M. Rogati (2010). Modeling relationship strength in online social networks. In *WWW '10 Proceedings of the 19th International Conference on World Wide Web*, pp. 981–990.

Zheleva, E., H. Sharara, and L. Getoor (2009). Co-evolution of social and affiliation networks. In *KDD '09 Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.